

From: vanthof (vant)
Sent: Wednesday, May 31, 1995 10:28 PM
To: Tim B. Robinson
Cc: vanthof (Dave Van't Hof)
Subject: Re: euterpe lvs with new padding lvs finished

Tim B. Robinson writes:

>
>Dave, what's the current DRC status? I was under the impression that
>under the current rules, and modulo any problems that the new pads
>could create at the top level, we are clean. However, at the pollux
>meeting this afternoon, mark said we had not run any euterpe DRC's for
>more than a month and the current status is unknown.

>
>Tim
>

Yes, it has probably been a month since any official drc's have been run.
Mostly because during the past month, nothing has been stable long enough to make drc's
reliable, and there was always the threat of more changes.

The database is probably okay. The pad edits were drc'd stand alone as were most other
drc edits that I know of. Another reason why drc's have been delayed is because of the
huge metal back-end synthesis requirements.
If we don't include the synthesis as part of a normal drc run, we will have many many
surprises at tapeout. The kicker is that the backend synthesis has only recently
stabilized. I'm in the process of coding up the rules, but these are incredibly complex
and every time I think I have a solution for even one layer, some unforeseen interaction
occurs which causes a reset on other layers. It's a very frustrating experience.

I talked with Mark about the drc flow this afternoon and to expedite things a bit, I'm
going to introduce the new rules into the existing drc flow tonight and run the fullchip
with that flow. This does expose us to errors which will only occur during the synthesis
at tapeout, but we can at least make some progress.

I only wish there were 48 hours in a day cuz I could sure use them.

Dave

--
Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100
"Don't blame me! I didn't vote for him"

From: doi (Derek Iversen)
Sent: Wednesday, May 31, 1995 8:44 PM
To: guarino; gmo; gregg; claseman; lisa; jeffm; lisar
Cc: hestia
Subject: Software Bringup Meeting Minutes - May 31, 1995

Software Bringup Meeting

May 31, 1995

Next Meeting: June 7 at 2:00 pm.

Attendees: guarino, gmo, doi, gregg, claseman, lisa, jeffm, lisar

New Action Items

Item: Need verify/ukernel Makefile modified to build other tests
Who: doi
Status: New

Review of Action Items

Item: Modify startup code in stb/stand to read the cerberus node number
Who: gmo
Status: Pending

05/03 No new progress

Item: Plan for testing remote debugging environment
Who: everyone
Status: Pending

04/13 This will be discussed at the next meeting.
04/20 Short discussion on the pieces required and if we want to have a standalone version of the hostio software.
Wally has started on this already.
04/26 Any work with snoopy/dram models will begin after euterpe tapeout when jeff has more time.
05/03 Wayne F. expects to have a prototype ISA card ready for testing tomorrow. Gmo is in the process of writing code to test the card.
05/10 Gmo has written a test program. Gmo and Wayne need to get together and try it on the board.
05/24 Wally has added more connectivity between the ukernel and the debugger (ability to interrupt kernel via the cerberus 'forced interrupt' bit).
Gmo is ready with a test when Wayne has the board ready.

Item: ukernel needs to detect machine checks and 'do the right thing'
Who: lisa
Status: Pending

04/13 No new progress.
04/20 On list of things to do. Lower priority.

Item: Create performance test plan

Who: claseman

Status: [11/30] No progress as focus is on functionality.

11/30 We continue to run tests to help us compare terp vs hardware performance.

We still need to put together the actual performance tests that need to be run on the hardware.

05/10 Back to life. Tim has checked in more tests. We are going to use the tests written for hermes and cerberus read accesses in the 'How to debug Euterpe' presentation on monday. We will be looking for the time it takes from the instruction issue, entry to nb, and completion of access.

05/24 Tim is going to send lisar a list of the test names that he intends to write so they can be incorporated into the test template.

05/31 Jeffm is going to analyze approx one test per day (with Tim looking over his shoulder) to determine/verify the actual performance numbers so they can be incorporated into the software simulator.

Suspended Items

Item: Unsnap code

Who: lisar, jeffm

Status: Suspended

02/15 The issue of restarting the hardware from an IKOS dump was discussed and the need for an architectural snap/unsnap facility was questioned.

Since the meeting it has been re-discovered (jeffm wasn't there to remind us of an earlier decision) that we are planning on loading architectural state into an IKOS simulation and not from a total IKOS logic dump.

We also determined that when it came time to run some of the larger tests (real-time benchmark) we would need the capability to start an IKOS simulation from an architectural dump anyhow.

03/01 For the short term we are going to focus on a simpler approach for loading and running DVTs, the kernel, and kernel tests. This item will likely come back in April.

05/10 Back to life. Does the IKOS support RAM dump?

05/24 This item was suspended again. There are no resources allocated to this item at this time.

Item: When do we have a full calliope simulation available (IKOS)?

Who: lisar

Status: Suspended.

04/26 This topic was raised as we talked about when the Snap/Unsnap item should be brought back from the Suspended list.

05/10 Lisar was not able to attend the meeting.

05/24 Suspended. This is related to the snap/unsnap item (above).

Completed Items

None.

Terp Feature Status

inprog o Ifetch protection granularity

- performance vrs accuracy tradeoff
- inprog o Fetch instructions as octlets
- inprog o Accuracy wrt HW simulator(s?)
- o Better latency model for Calliope accesses
- o Implement hardware configuration through Cerberus regs (SDRAM parameters, dram enable?)
- o Checkpoints/Snapshots
- inprog o ability for terp to load hermes sections
- lisar would like this functionality added
- inprog o remove stbio from hwterp.

Performance Test Status

Status	Test Description	Verified	
		H/W	S/W
ran	store/load to dram	No	No
ran	store/load to hermes	No	No
ran	store/load to cerberus	No	No
ran	load from rom	No	No
ran	icachemiss	No	No
ran	dcachemiss	No	No
	load ltlb entry (write+read)	No	No
	load gtlb entry (write+read)	No	No
	NB overflow	No	No
	generate an interrupt	No	No
	return from interrupt	No	No
	mult cys trying to take exceptions at the same time	No	No
	predicted branch	No	No
	unpredicted branch	No	No
inprog	One instruction from each instruction class:	No	No
	arith store eshift imul64		
	branch swap epop idiv		
	branchx stored imul4 gfmul		
	load eset imul8 bgate		
	loadx eshift imul16 atomic ops		
	nbload eshiftx imul32		
inprog	Inner loop sequences:	No	No
	- cable_in main handler-- inner loop in		
	- EQ_UPDATE_WEIGHTS() (khp)		
	- IDCT code (fur)		
	- pseudo-Huffman decode loop (fur)		
	- a reconstruction routine from macroblock.c (fur)		
	- NTSC encode loop (fur)		
	- rs_compute_syndrome() (fur)		

Test Status and General Discussion

Lisa has build and is running a Euterpe/Mmemo simulator (Ikos) that we will run someof the hermes tests that jeffm has written (with some modifications to initialize mmemo added) and then we can start thinking about running some of the oc_mem tests.

It was discovered that caches interleave didn't work. A fix for the HW has been released this morning.

We need to modify the ukernel's taskExit so that tests running on the HW simulators can tell when the test finishes.

Work will begin to prepare for building and running Unix on the HW simulators. Gmo has determined that it will run for about 6 million cycles before doing a 'probe' (required use of hostio and this doesn't work yet on the HW simulators). We can fake this

part out and then run for quite a while longer before hitting a point we can't fake out a hostio requirement.

Derek and Wally will look at building an OSF kernel.

From: tbr
Sent: Wednesday, May 31, 1995 8:22 PM
To: hessam (Hessam Mohajeri)
Cc: graham; gmo; abbott; dbulfer
Subject: Telephone Modem

Hessam Mohajeri wrote (on Wed May 31):

Hi, Tim

I am not quite sure what we need to do as far as the modem concerns. Do we want to build the function on the main board using discrete components or we would like to have an AT-bus type interface.

It has to be an option. That could be

- a. an option on the main board, unpopulated when not needed
- b. two versions of the board
- c. an optional module with a connector

I strongly favor c. which is also consistent with what we've been hearing as the preference in meetings with @home. Later in high volume when a majority of systems have an RF uplink and the function is not required we can choose to depopulate the connector on the main board.

A major consideration would seem to be that if we can buy in a complete module at a competitive price, it would be pre-approved for connection to the phone network, eliminating many of the regulatory hurdles. Obviously it would have to be something in volume production which is priced aggressively, though we would not be expecting to count it in the BOM for the basic unit. Indeed we may not end up handling it at all if it is bought directly by the MSOs.

The first suggestion we came up with was a PC104 module. That appeared to have the dual advantages of small form factor and simple interface (logically equivalent to ISA, but a different connector). However, it seems these are more targetted to industrial applications and so not driven to low cost and are a generation behind in performance.

More mainstream are either PCMCIA or vanilla ISA cards. Both of these are available retail qty 1 in Fry's for order \$30, so clearly there is the possibility of getting them at an acceptable price. PCMCIA is much smaller, but has a more complex interface needing a greater pin count than ISA. In considering what bus interface we should put directly on our chip, ISA is a strong contender because:

- a. it's simple
- b. low signal count
- c. may be easy to share the same interface with flash ROM
- d. other chips are available to interface directly with no glue
(we are looking at at least one ethernet chip with this interface)
- e. Form factor may not be a problem given the box size will likely be driven by large heatsink dimensions, so a short ISA card might be no problem to accommodate.
- e. PCMCIA interface would likely be a problem for pincount and logical complexity. (When we considered it as a direct interface on euterpe, we ended up thinking we would need to use an external chip to map from the lower pincount flash interface, and to provide the address mapping registers it requires).

Now, given the desire to get to very low cost in this product we need to be flexible in what we integrate into our chip so that we can work with minimal (preferably zero) external glue chips with the lowest cost external devices. So, I would prefer we do the homework to find what has the potential to be the overall lowest cost in volume, taking

into consideration all the system level factors (eg connectors, power/cooling requirement, die area needed for the interface, rather than start out with fixed ideas about what our interface should be.

If anyone has any other suggestions as to low cost ready made modules, let's investigate them too.

Tim

From: tbr
Sent: Wednesday, May 31, 1995 7:30 PM
To: vanthof (vant)
Subject: euterpe lvs with new padring lvs finished

vant wrote (on Wed May 31):

The euterpe lvs with the new fixed pad ring came back clean. There was one problem with the lvs and it really needs to be redone (which I'll startup right away). I did not realize that two new cells were created as part of the new pad changes so these cells were not included in the lvs run. They dealt with the corner pads and crack/seal ring. Lookin good so far.

Dave, what's the current DRC status? I was under the impression that under the current rules, and modulo any problems that the new pads could create at the top level, we are clean. However, at the pollux meeting this afternoon, mark said we had not run any euterpe DRC's for more than a month and the current status is unknown.

Tim

From: paulp (Paul Poenisch)
Sent: Wednesday, May 31, 1995 5:57 PM
To: cadettes; al (Albert Matthews); anh (Anh Ngo); calliope; graham (Graham Y. Mostyn); abbott (Curtis Abbott); mouss (John Moussouris)
Subject: Calliope1.1

We are currently processing Calliope1 desive through the fab. The reticle set that we have for this devices works well enough that we can, usually, process the wafers all the way through the line. However it has become clear in the past few weeks that the problems that this reticle set has will likely prevent us from ever yielding a part with it. We may get a few partially functional parts if we run enough lots but there will be very few and they will not be very functional.

The problems with this set of reticles are the same ones that we have been working on fixing on Euterpe. They primarily involve the I/O ring and associated structures (seal ring, power distribution, corner cells, etc.).

If we want to get Calliope's that work well enough for circuit and software debugging we need to have these identified problems fixed. I believe that the cells that have been (or are being) modified for Euterpe cover most of the cells that would need changing on Calliope. Therefore it is the feeling of the fab organization that we need a "Calliope1.1" reticle set. This reticle set should have a minimum of changes related to circuit issues (read none) to allow us to turn it as quickly as possible. It would need to be a complete set due to the changes that were made in some of the cells.

Obviously a new Calliope tape out will need to be scheduled in with the other tape outs that are coming up. A tentative priority might be Euterpe, Pollux, Castor, Calliope1.1, Mnemo (as discussed in a meeting with Geert, Test and CAD today). Clearly this is an issue that will need discussion based on the needs for a working Calliope.

If you have questions, comments or over ripe fruit involving this issue please let me know.

Paul.

From: hopper (Mark Hofmann)
Sent: Wednesday, May 31, 1995 9:53 AM
To: vant
Cc: tbr (Tim B. Robinson); lisar (Lisa Robinson); vo (Tom Vo); geert (Geert Rosseel); vanthof (Dave Van't Hof); wampler (Kurt Wampler); tom (Tom Laidig)
Subject: Re: euterpe lvs with new padding lvs finished

vant writes:

The euterpe lvs with the new fixed pad ring came back clean. There was one problem with the lvs and it really needs to be redone (which I'll startup right away). I did not realize that two new cells were created as part of the new pad changes so these cells were not included in the lvs run. They dealt with the corner pads and crack/seal ring. Lookin good so far.

Dave

Okay. This is good, very good. Let's re-start that LVS and keep hold of this design. Once the _last_ pad edits are done and locked down, let's LVS again and start a tapeout of Euterpe (first the bottom layers, and then the top as soon as the flow is ready)

-hopper

From: hopper (Mark Hofmann)
Sent: Wednesday, May 31, 1995 8:06 AM
To: cadettes
Cc: tbr (Tim B. Robinson); geert (Geert Rosseel)
Subject: tapeout work

As a result of a CAD pow-wow today, here are the latest estimates of time to Euterpe tapeout. Note that, as before, these numbers assume that the DRC flow is stable and that no re-work will be needed. So it should be taken as a minimum guideline.

Weeks	Task
0	SDEC synthesis
4	Metal synthesis [hole perforation, notch filling, via grow, waffle, make fat metals, update fracture and DRC flows]
0	DRC flow maintenance
1	Frame cell mods
1	scribe fill cells
0.5	scribe/die interface
0.5	copyright/dev ids
1	Eu hand route
2x6C	Eu DRC [2 turns @ 1 week on 6 machines]
2x6C	Po DRC [2 turns @ 1 week on 6 machines]
2x0.75C	Eu LVS [2 turns @ 0.75 weeks each]
2x0.75C	Po LVS [2 turns @ 0.75 weeks each]
0	Fix Eu LVS
5-6x2C	fracture Eu [5-6 weeks on 2 machines]
5-6x2C	fracture Po [5-6 weeks on 2 machines]
2	Vlsimm enhancements
0	Eu/tools snap
0	Eu csyn
1	Eu celltest
0.5	Eu Space xformer
(3)	manual SDEC fix [Now a "guideline", not a rule]
0	mobieclium noxistors
0.5	sofagen/make Po (+ Rich/Johnny)
3	Rev Po to current Mobi rules
0.5	fracture Eu space xformer
0.5C	fracture Eu space xformer
2	gen space xformer frame
0.5	minimum bias improvements [all metals +]
1	check post-fracture DRCs
---	-----

19 Man weeks / 4 people fulltime => 5 weeks
 51.5 CPU weeks / 4 machines in parallel => 13 weeks

Assume that CPU cycles prior to fracture can be overlapped with man weeks for about a half overlap or 6.5 weeks.
 Therefore the first mask can be created in about 5 weeks (new flow) and the last mask will be finished in 5 + 6.5 = 11.5 weeks

To tapeout just the lower layers of Euterpe (010 - 140 and 160) if we use the current (old) DRC flow would take us about 3 weeks:

1 week to fix pads, edit crack protect ring, DRC, LVS
 2 days to check interaction between space transformer pads and baspelate
 1 day to fix mobieclium noxistors + run shorts test
 1 week to work on scribe fill pattern issues. [in parallel with layer DRC, assumes etest patterns stable]
 1 week to run DRCs (6 machines) [assume clean on first run!]
 2 days to fracture

Depending on amount of overlap, this is about 2.5-3 weeks to generate tapes for the lower layers.

-hopper

From: wingard (Drew Wingard)
Sent: Wednesday, May 31, 1995 2:23 AM
To: cronus
Subject: Summary of Cronus/Atlas Status Meeting, May 26

Hi,

Here is a summary of last week's Cronus/Atlas Status meeting

Attendees: ong, brianl, wampler, hopper, serena, mudge, tau, kgk,
paulb, steve, tbr, lisar, warren, wingard

Note: Lines beginning with ">" are from Geert's message (of 5/9) reporting the results of the meeting on May 5th.

> 1. Baseplate
>
> Current status : We have a gards-compiled cronus baseplate
> which can be used for routing experiments,
>
> Plan : By May 31, we want to have a final baseplate.
>
> Action items : * Decide on final die size : Drew
5/19 - Makefile and tools nearly in place to support this
activity. Should make June 1 deadline, but in
some danger of slipping.
5/30 - Tools now working. Floorplanning finally begins.
Won't make June 1. Hope to make June 5.

> * Padring assignment : Warren
5/19 - On hold pending floorplanning.
5/26 - On hold pending floorplanning.

> * Floorplan, custom block
> placements : Warren
5/19 - On hold pending die size work. Power interface
(hemming) cell work is ongoing. First version of
clock mast cut-outs released.
5/26 - On hold pending floorplanning.

> * Clock Tree generation : Kurt, Bill
5/19 - In progress. Bill has communicated plans for a
single-mast clock system to Kurt. Kurt indicates
on-schedule completion is very likely.
5/26 - Substantial progress. Ready for clock buffer
cell (currently in layout).

> * Build dummy cells for
> TTL I/O and : B.P
> IOBYTE : Dane
5/19 - No progress to report.
5/26 - No progress to report.

> 2. Custom blocks
>
> Current status : GTLB done
> Register file : end of this week
> Caches : layout at top-level
> Tags : layout has not started yet
> NB : layout has just started

> TTL I/O : not designed yet
> IOBYTE : not started yet
>
> Plan : finish all custom blocks by end of May
> except TTL I/O and IOBYTE
>
> Action Items : * Register File : B.P, Mike, Orlando
5/19 - Polygons finished on 5/19. Largely DRC/LVS clean.
Top-level should be LVS clean by 5/24.
5/26 - No report, but appears to be clean.
>
> * Caches : Bruce, Efeliass, Dtacmo
5/19 - Laying out 'last' two cells. Final size is known.
5/26 - No report, but doesn't appear finished.
>
> * Tags : Bruce, Efeliass, Dtacmo
5/19 - No progress, but can use only the cache leaf
cells and therefore be only top-level assembly.
5/26 - No report, but some blocks checked in.
>
> * NB : Vikki, B.P.
5/19 - Polygons should be finished by about 5/24.
5/26 - No report, but doesn't appear finished.
>
> * TTL I/O : B.P., ??
5/19 - Design goals specified. In design. Some early
layout work has begun?
5/30 - Circuit design "complete"
>
> * IOBYTE layout can start by end of next week.
> We'll need schematics by then : Dane, Bill
5/19 - No progress was reported.

5/19 Action Items: * IOBYTE is apparently slipping badly. It will not
make the June 1 deadline. Drew to find out what we
can do to make better progress.
5/26 - Bill, Dane, Tbr, and Drew met to better understand
issues. Analog portions in layout. Digital
delay line and ECL-CMOS converter to await Bill's
return.

* XLU: assumption was that we would try to build XLU
datapath out of Atlas std cells (with 'custom'
placement and hand route). Tbr to ask Karzes about
modifying generation program to output std cells.
Drew to consider adding mux16 to support Stage 2.
5/26 - Karzes indicates that modifying generation
program would not be difficult if required.
Kleanthes to drive Cronus XLU design, starting
with existing std. cells.

* CERPOKGEN: Power OK circuit needs to be designed.
Since the caches will be a client, consider using
circuit that Bruce had started designing.
5/26 - No progress to report

> 3. Atlas database

>
> Current status : Database builds from toplevel. An initial version
> of about everything is there.
5/19 - Atlas database builds from scratch as of 5/18

>
> Plan : Have a complete and accurate Atlas database by the end of May
> In building Cronus sub-blocks, everybody should point to
> /u/chip/atlas and not local versions. That will catch
> database problems early.

```

>
> Action Items : * Review XL, XS cells
>                 timing, layout, cap models : Warren
>
> 5/19 - All cells have passed corner analysis. Some
>        had to be resized for better margin. Latches
>        and muxes have been performance tuned. Other
>        cells ongoing.
>
> 5/26 - Latch timing found to be more difficult than
>        expected. Analysis continues.
>
>
> * Finish timing numbers : Fred
>
> 5/19 - No progress to report. Hopper to help with
>        capacitance and input loading work.
>
> 5/26 - Some progress on flip-flop timing. Unlikely
>        to make June 1. Early LPE flow (for capacitance
>        extraction) installed.
>
>
> * Complete verilog libraries : Brianl
>
> 5/19 - Basic libraries installed. Still need models
>        for Cerberus and 'SC' cells.
>
> 5/26 - Cerberus and 'SC' simulation models installed.
>        Plan to rely on Synopsys to map most of these
>        into std. cells. More automated (schematic-
>        based) extraction of models in progress.
>
>
> * Complete Synopsys libraries : Brianl
>
> 5/19 - Basic libraries installed. Severe performance
>        degradation with latest version of Synopsys
>        noticed, but can still use prior version until
>        fix/workaround is provided.
>
> 5/30 - Modifications underway to use 'proper' timing
>        numbers from atlas/leafchar.
>
>
> 5/19 Action Items: * Enhance Verilog libraries : Brianl and Drew
>                    Tbr has provided a list of 'missing' cells
>                    that prevent him from compiling Cronus for
>                    simulation at the top level. We need to get
>                    the truly needed ones installed ASAP.
>
> 5/26 - Cerberus and 'SC' simulation models installed.
>
>
> * Enhance Gate Family : Warren and Drew
>
> We need XOR and MAJ gates, and perhaps an
> asynchronously-loadable latch for Cerberus.
>
> 5/30 - In progress.
>
>
> 4. Cronus/verilog/bsrc, Makefile.rules, GARDS
>
>
> Current status : Initial versions of Makefile.rules exist. We can
>                  build a sub-block using the atlas database on
>                  the checked in baseplate model.
>
>
> Plan : Have a "good" Makefile.rules that works well enough that
>        other people can start mapping Euterpe blocks by the end of
>        May
>
>        Implement " put the top-level together" strategy in the
>        equivalent of euterpe/verilog/bsrc/Makefile.tst
>
>
> Action Items : * Makefile.rules : Brianl
>
> 5/19 - "Better" version released. Ongoing.
>
>
> * GARDS model additions to deal
>
> with power and clock obs : Tau, Kurt
>
> 5/19 - Completed and tested - works great!
>
>
> * Pim2pif upgrade : Hopper
>
> 5/19 - Completed.

```

> * Take euterpe snapshot : Tbr
 5/19 - Completed. Cronus.V released.
 Tbr has been cutting out Euterpe blocks that Cronus will not have, and bringing over behavioral simulation models of the custom blocks. Billz has swapped in a behavioral NB buffer description. Dickson has been modifying the Cerberus register map to remove bits associated with unused functionality.
 5/26 - Cronus has run a few cycles in simulation!

> * Floorplanning tools : Drew
 5/19 - Makefile and tools nearly in place to support this activity. Should make June 1 deadline, but in some danger of slipping.
 5/30 - Largely done, but probably not in time.

> * verilog/bsrc/Makefile : Drew
 5/19 - See above. Rest of this activity is largely in the hands of Tbr and Brianl.

> 5. "Implementation of mapping"

> No plans for this month
 5/19 - Brianl has been 'test' mapping the top-level Euterpe blocks to identify mapping and P&R issues. This should also provide valuable feedback in setting the die size.
 5/26 - Brianl has released some early mapping results that have encouraging area ratios (only 2:1 over Euterpe). However, these results do not include and P&R information. More info this week.

> 6. Packaging

> Current status : We have "a plan" (read all about it in Mosaic)
 > Plan : By the end of May, make decisions on all aspects of "the plan"
 > Action Items : * Call a set of meetings during May to turn our plan into a set of decisions : Geert
 5/19 - Met on 5/15 and decided on 70mm TAB frame. Trancy investigating having someone else do ILB/OLB work, but lead times are about as bad as doing it ourselves. Drew spoke with MMS and got clarification on module spacing requirements. Trancy needs firm substrate size to proceed much further. Drew to provide that this week.
 5/26 - In trouble here. Need to order *NOW* to have packaged parts shortly after wafer delivery. MMS wants us to deal with wafer bumping company (Aptos) directly. ILB after chip & capacitor attachment looks tricky/risky.

5/26 Action Items: * To enable parallel progress on packaging/testing setup, we will 'tape out' the pad ring of Cronus on 6/16. We have also set this as the date to begin maintaining a Cronus 'snapshot' environment.

> 7. Test

> Current Status : Mudge & Mark Warren have taken the responsibility for Cronus test, both die sort and test of packaged parts,
 > Plan : Same as packaging : by the end of May, we want to decide

> on our plan so that only implementation issues are left.
>
> Action Items : Johnny and Mark own this, so they can decide.
5/19 - No progress was reported.
5/26 - MMS is concerned about non-uniformity of internal
pads with respect to Known Good Die test solution.
Drew promised a representative die plot by 6/2.

Verification:

5/19 Action Items: * Need to csyn/celltest custom blocks. Major missing
piece at this point is the Verilog. We could use
some logic/verification help on this one.
5/26 - Lisar presented a schedule for Cronus behavioral
and gate-level verification. Csyn/celltest
should be a subset of Euterpe, and thus simpler.

* Atlas-based designs have latches, so we need to do
phase checking. Hopper to work on getting 'gloss'
running for Atlas.
5/26 - Gloss has been successfully run on an Atlas block

Regards,
Drew

From: wampler (Kurt Wampler)
Sent: Tuesday, May 30, 1995 3:36 PM
To: nebabiek@silvix.com
Cc: brianl; geert; hopper; tbr; tom; wampler; wingard
Subject: Re: Maze routing problem

Hello, Nebabie -

I dropped off an envelope for you this afternoon (the receptionist signed for it), containing examples of the GAROUT and RROUTE problems that I described to you in my e-mail messages of May 5, and May 15, respectively.

The GAROUT example appears to be a fairly serious bug, and we would like to flag it as "high priority".

The RROUTE bug is less serious; I believe I figured out what tickles this particular bug, and the example on the tape demonstrates this. Now that we know what causes it, I believe we can work around it, so I would flag this one as "medium priority" -- I would like to have an SVR developer examine the source code and make sure this one isn't more sinister than it appears on the surface -- having it fixed in the next release would be fine. (There is some possibility that the root cause of this bug may be related to our earlier unresolved TAR #2004.)

Please give me a call when you get back from vacation tomorrow and confirm for me that you received the tape and that it reads in successfully.

- Kurt

From: tbr
Sent: Tuesday, May 30, 1995 11:11 AM
To: dbulfer (David Bulfer)
Cc: tbe@microunity.com
Subject: Re: my wish list

David Bulfer wrote (on Tue May 30):

>
> Tom, I don't recall following up on this. Was anything done? The
> only question I would have is to be sure 4mm DAT is the right choice.
> Most of the tapes we send out go in Exabyte 8mm format. Either way,
> seems like it would be desirable for you to have it at your desk.
>

Sorry, I think that I let this fall through the crack. I spoke with Frank about it and we were going to get you one. I can't remember if Frank was waiting on me for something or if I was supposed to talk to Ken. My Alzheimer's must be kicking up again.

I'll check into it.

I'm assuming gearhead is the SGI. If so talk to ken.

Tim

From: hopper (Mark Hofmann)
Sent: Tuesday, May 30, 1995 1:41 AM
To: Tim B. Robinson
Subject: Re: More detailed estimate of CAD tasks to tapeout Euterpe

Tim B. Robinson writes:

Mark where do we stand in progress relative to this list?

Tim

Mark Hofmann wrote (on Thu May 11):

The CAD folks met today to try to quantify the amount of work remaining between now and a full Euterpe maskout. The numbers below assume that the DRC flow is stable and that no re-work will be needed.

So it should be taken as a minimum guideline.
[snip]

Hi Tim,

Let me get back to you on this. The rough answer is: 1. No work has been done by the CAD folks on Pollux. (it now appears it would go on a separate reticle set). 2. Fwo/Drew are still trying to answer one final Celltest question on Euterpe (this was the long running Tcl job) but Euterpe looks good from a Celltest standpoint, I think. 3. Dave/Tom have taken a new approach from what was list on 11 May to come up with a single "unified flow" that handles waffle, synth, fracture and DRC as one so the tasks will come out a little differently. I will try to enumerate a new list from where we are now to Euterpe tapeout.

-hopper

From: tbr
Sent: Sunday, May 28, 1995 11:17 PM
To: Tom Eich
Cc: dbuffer
Subject: my wish list

Tom Eich wrote (on Sun May 14):

Hi,

I realize that this is not something that I can't live without, but it would= be quite helpful if I had a 4mm DAT on gearhead. On the sole occasion when= I've required a 4mm DAT, I've needed help from Kurt W. (who has one of the= company's two 4mm drives in his office) in writing the files to tape. = The plan for the Pandora fabs is to dump many designs for sheet metal and= plastic parts onto tape and ship via vendor truck. I anticipate that this= will happen quite a bit in the next several weeks.

Rather than having to run down to Kurt W.'s or Khp's office to load the= tape, and possibly interrupt one of them in the process, I would like to= take care of everything in one place. I believe that having a DAT attached= to gearhead would reduce the time and number of resources required to get= these tapes written. The drive could be centrally located in the boxer's= print room, if having it in my cube is a problem. Please let me know if I= can plan on having such an asset.

Tom, I don't recall following up on this. Was anything done? The only question I would have is to be sure 4mm DAT is the right choice. Most of the tapes we send out go in Exabyte 8mm format. Either way, seems like it would be desirable for you to have it at your desk.

Tim

From: tbr
Sent: Sunday, May 28, 1995 10:56 PM
To: hopper (Mark Hofmann)
Subject: More detailed estimate of CAD tasks to tapeout Euterpe

Mark where do we stand in progress relative to this list?

Tim

Mark Hofmann wrote (on Thu May 11):

Hi,

The CAD folks met today to try to quantify the amount of work remaining between now and a full Euterpe maskout. The numbers below assume that the DRC flow is stable and that no re-wrok will be needed. So it should be taken as a minimum guideline.

Weeks	Task
0	SDEC synthesis
2	Metal Waffle
2	Metal synthesis
0	DRC flow maintenance
1.5	Frame cell mods
1.5	scribe fill cells
1	scribe/die interface
1	copyright/dev ids
1	Eu hand route
2x6C	Eu DRC
2x6C	Po DRC
1.5C	Eu LVS
1.5C	Po LVS
1	Fix Eu LVS
2	mod. Mobi fracture flow
5-6x2C	fracture Eu
5-6x2C	fracture Po
3	Vlsimm enhancements
0	Eu/tools snap
0	Eu csyn
1	Eu celltest
0.5	Eu Space xformer
0.5	via enlargement check
3	manual SDEC fix
0	mobieclium noxisters
0.5	sofagen/make Po (+ Rich/Johnny)
8	Rev Po to current Mobi rules
0.5	fracture Eu space xformer
2	gen space xformer frame
---	-----

32 Man weeks / 4 people fulltime => 8 weeks
51 CPU weeks / 4 machines in parallel => 13 weeks

Assume that CPU cycles prior to fracture can be overlapped with man weeks for about a half overlap or 6.5 weeks.
Therefore the first mask can be created in about 8 weeks and the last mask will be finished in $8 + 6.5 = 14.5$ weeks

-hopper

From: tbr
Sent: Sunday, May 28, 1995 10:56 PM
To: hopper (Mark Hofmann)
Subject: More detailed estimate of CAD tasks to tapeout Euterpe

Mark where do we stand in progress relative to this list?

Tim

Mark Hofmann wrote (on Thu May 11):

Hi,

The CAD folks met today to try to quantify the amount of work remaining between now and a full Euterpe maskout. The numbers below assume that the DRC flow is stable and that no re-wrok will be needed. So it should be taken as a minimum guideline.

Weeks	Task
0	SDEC synthesis
2	Metal Waffle
2	Metal synthesis
0	DRC flow maintenance
1.5	Frame cell mods
1.5	scribe fill cells
1	scribe/die interface
1	copyright/dev ids
1	Eu hand route
2x6C	Eu DRC
2x6C	Po DRC
1.5C	Eu LVS
1.5C	Po LVS
1	Fix Eu LVS
2	mod. Mobi fracture flow
5-6x2C	fracture Eu
5-6x2C	fracture Po
3	Vlsimm enhancements
0	Eu/tools snap
0	Eu csyn
1	Eu celltest
0.5	Eu Space xformer
0.5	via enlargement check
3	manual SDEC fix
0	mobieclium noxisters
0.5	sofagen/make Po (+ Rich/Johnny)
8	Rev Po to current Mobi rules
0.5	fracture Eu space xformer
2	gen space xformer frame
---	-----

32 Man weeks / 4 people fulltime => 8 weeks
51 CPU weeks / 4 machines in parallel => 13 weeks

Assume that CPU cycles prior to fracture can be overlapped with man weeks for about a half overlap or 6.5 weeks.
Therefore the first mask can be created in about 8 weeks and the last mask will be finished in $8 + 6.5 = 14.5$ weeks

-hopper

From: hopper (Mark Hofmann)
Sent: Sunday, May 28, 1995 4:22 AM
To: vant; wampler (Kurt Wampler); tom (Tom Laidig); albers (Daniel Albers)
Cc: tbr (Tim B. Robinson); paulp (Paul Poenisch); geert (Geert Rosseel)
Subject: Re: SEM results

Kurt Wampler writes:

[snip]

BTW - Fung was showing me some SEM photos of excellent metal results Friday evening; he said they changed the metal etch step and got dramatically better results -- he seemed very encouraged. The metal photos were of Castor/Pollux/Calliope0 too (which didn't have the minXmin special bias etc.) and all of the minXmins were there -- everything looked picture perfect. This raises my hopes that with all of the enhancements applied to the Euterpe masks, we might actually see some measurable yield.

Sounds promising.

I think next week we may have more discussion about pushing out the first N masks ($10 \leq N \leq 14$) so that we can have masks ready for the fab to work on now. Perhaps Al will want to make some "final" corrections, I don't know. We should get together the interested parties on Tuesday and see what we need to do to actually make mask tapes of the bottom layers.

thanks,
-hopper

From: vanthof (vant)
Sent: Saturday, May 27, 1995 2:55 PM
To: mikew (Mike Wageman)
Cc: vanthof (Dave Van't Hof); geert (Geert Rosseel)
Subject: PAD CELLS CHECKED IN

Mike,

I checked in your pad cell edits today. From my understanding, you had completed the edits and were simply waiting for the 'okay' to check them in. The fullchip lvs finished clean so we would like to run another check with the latest pad cells we want to tape out, AND since you are goofing off and taking a weekend away from work, someone had to check in these cells :-)

Thanks,
Dave

--

Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100
"Don't blame me! I didn't vote for him"

From: hopper (Mark Hofmann)
Sent: Thursday, May 25, 1995 1:47 PM
To: Tim B. Robinson
Subject: Re: cpu cycles for CAD tapeout

Tim B. Robinson writes:

I spent quite a while talking to mouss about some of this and I had to say I had no answer as to why breaking it up could not help with the memory issues. His point being that most of the work is inherently local.

It's not that Mouss is on the wrong track, it's that the issue is more complicated when you actually get down to the implementation details.

Here's mail from Tom:

Tom Laidig [tau] writes:

If we do strips, we obviously have to have a certain amount of overlap to assure that the synthesized geometry on both sides of a boundary match. Dave and I spent a bit of time trying to figure out how big that overlap would need to be. Clearly, we need at least 4.5u on each side of the boundary (ie -- 9u overlap) to be sure that both sides agree on whether some straddling piece of metal gets perforated. Dave and I believe that some effects can propagate horizontally as we move up the layer stack, but we got lost trying to guess how far this could go. It is possible that we'd need to add about 4u for each half-layer in either the bottom-up or top-down part of the flow, so we might be talking about something like a 30u border around each strip (60u overlap). I just don't have much confidence that we really know how big a border is needed.

Further, the notch filling step cannot be guaranteed if the chip is divided into chunks in any way, regardless of border width. A notch or hole that spans the whole chip would only be seen to be a notch at one end (two end for a hole), and would be seen as a legal gap in the middles strips. I don't know how long a notch might be, but the construction of the sofa means we have Vdd notches that span the entire distance between clock spars. Note that this problems is not caused by the fact that we define the rule as a notch rule instead of a hole rule. Indeed, the current formulation is better, since we can at least check for illegal notches using strips, whereas this would be impossible for a hole-checking formulation.

I'm afraid that the more I think about it, the less I like the striping idea. In addition to the problematic nature of the border width, there's the issue of intermediate disk storage space for all the fully-synthesized layer strips, the fact that each thread would need to flatten all layers (cpu expense; require >=128MB ram to avoid thrashing), and the fact that we'd waste cpu time writing all this temporary data and then reading it back in for the final flat fracture thread (unless we figure out how to fracture in strips and splice the fractured results together -- I guess we could just tape out a set of ten or so 'chips' that the mask shop can place adjacent to each other on the reticle... ugh!). Overall, it just seems like a big, never-ending can of worms (every time we tape out, we'll have to review the striping plan) for less benefit than we would get by implementing intelligent parallelism in vlsimm (where we completely avoid duplicated flattening work, intermediate disk storage, and gratuitous read/write time). I'd frankly feel better about spending my time enhancing vlsimm than scratching my head about how to stripe each chip we tape out.

[End of Tom's mail]

Therefore, we thought of commissioning the sysadm group with two tasks:
1. find out specs of the Power Challenge: how many processors, how much memory, how much swap, how much memory per process, is anything 64-bit, how does it compare with our current challenge? 2. same questions of the DEC alpha. Both of these machines would speed up DRC, waffle, synthesis and fracture. In addition, a true 64-bit architecture might speed up LVS (which is not a bottleneck, at present).

We should do the research. However, I'm very worried about our experience to date with SGI and in particular their unwillingness for a long time to really acknowledge and address the problems. I also feel they have fallen well short of their promises. So we should be wary of anything that's not actually shipping.

I agree with what you say, but I think we could take advantage of a shared memory machine with fast processors. I did speak with Bob a bit today. He said his friends at the NAS project had one of the very early Power Challenge boxes. He should be able to get us good information on these machines.

As regards LVS not being a bottleneck, while we may not have much we can run in parallel right now, the runtime is definitely a major concern, which is why I have been asking frequently about ISS.

LVS could be a bottleneck in the future. Of the steps: waffle, synthesis, DRC, fracture, XOR check, LVS- LVS is the fastest at 3-4 days on a single processor. By comparison, DRC is 7 days across 6 machines. The problem with ISS is, unfortunately, that it is tuned for speed at the expense of memory. Dave ran out of VM on Godzilla when he tried his last Euterpe run. ISS runs DEC Alpha's. I think we might need to go to that architecture (and add lots of main RAM) to take advantage of ISS's performance.

Let's talk some on this and then get with Bob to discuss what to do.
Tim

Okay, good. I have kind of a full Friday- how about late morning or early afternoon?

-opper

From: tbr
Sent: Thursday, May 25, 1995 5:41 PM
To: hopper (Mark Hofmann)
Subject: cpu cycles for CAD tapeout

Mark Hofmann wrote (on Thu May 25):

Hi Tim,

The CAD folks have been pushing around the idea of adding hardware to speed up back end tapeout processing. The conclusion is that breaking a chip into strips is probably more difficult than it seems at first blush and doesn't address memory problems. What does seem promising is to make Vlsimm run in parallel. If tom added this (it is something he's thinking about) and we had a multi-processor machine with shared memory we might see a significant gain.

I spent quite a while talking to mouss about some of this and I had to say I had no answer as to why breaking it up could not help with the memory issues. His point being that most of the work is inherently local.

Therefore, we thought of commissioning the sysadm group with two tasks:
1. find out specs of the Power Challenge: how many processors, how much memory, how much swap, how much memory per process, is anything 64-bit, how does it compare with our current challenge? 2. same questions of the DEC alpha.
Both of these machines would speed up DRC, waffle, synthesis and fracture.
In addition, a true 64-bit architecture might speed up LVS (which is not a bottleneck, at present).

We should do the research. However, I'm very worried about our experience to date with SGI and in particular their unwillingness for a long time to really acknowledge and address the problems. I also feel they have fallen well short of their promises. So we should be wary of anything that's not actually shipping.

As regards LVS not being a bottleneck, while we may not have much we can run in parallel right now, the runtime is definitely a major concern, which is why I have been asking frequently about ISS.

Let's talk some on this and then get with Bob to discuss what to do.

Tim

From: tbr
Sent: Thursday, May 25, 1995 5:33 PM
To: lisar
Subject: forwarded message from Derek Iverson

Is *anything* getting done?

----- Start of forwarded message -----

Status: RO
X-VM-v5-Data: ([nil nil nil nil nil nil nil nil
["7635" "Thu" "25" "May" "1995" "10:03:29" "-0700" "Derek Iverson" "doi" "nil" "222"
"Software Bringup Meeting Minutes - May 24, 1995" "^From:" nil nil "5"]])
Return-Path: <doi>
Received: from polyhymnia.microunity.com by gaea.microunity.com (4.1/musel.3)
id AA23812; Thu, 25 May 95 10:03:30 PDT
Received: by polyhymnia.microunity.com (8.6.10/muse-sw.3)
id KAA07426; Thu, 25 May 1995 10:03:29 -0700
Message-Id: <199505251703.KAA07426@polyhymnia.microunity.com>
From: doi (Derek Iverson)
To: guarino, gmo, doi, gregg, claseman, lisa, jeffm
Cc: hestia
Subject: Software Bringup Meeting Minutes - May 24, 1995
Date: Thu, 25 May 1995 10:03:29 -0700

Software Bringup Meeting

May 24, 1995

Next Meeting: May 31 at 2:00 pm.

Attendees: guarino, gmo, doi, gregg, claseman, lisa, jeffm

New Action Items

None.

Review of Action Items

Item: Modify startup code in stb/stand to read the cerberus node number

Who: gmo

Status: Pending

05/03 No new progress

05/10 Ditto.

05/24 Ditto.

Item: Plan for testing remote debugging environment

Who: everyone

Status: Pending

04/13 This will be discussed at the next meeting.

04/20 Short discussion on the pieces required and if we want to have
a standalone version of the hostio software.
Wally has started on this already.

04/26 Any work with snoopy/dram models will begin after euterpe tapeout when jeff has more time.
 05/03 Wayne F. expects to have a prototype ISA card ready for testing tomorrow. Gmo is in the process of writing code to test the card.
 05/10 Gmo has written a test program. Gmo and Wayne need to get together and try it on the board.
 05/24 Wally has added more connectivity between the ukernel and the debugger (ability to interrupt kernel via the cerberus 'forced interrupt' bit).
 Gmo is ready with a test when Wayne has the board ready.

Item: ukernel needs to detect machine checks and 'do the right thing'
 Who: lisa
 Status: Pending

04/13 No new progress.
 04/20 On list of things to do. Lower priority.
 04/26 Ditto.
 05/03 Ditto.
 05/10 Ditto.

Item: Create performance test plan
 Who: claseman
 Status: [11/30] No progress as focus is on functionality.

11/30 We continue to run tests to help us compare terp vs hardware performance.
 We still need to put together the actual performance tests that need to be run on the hardware.
 05/10 Back to life. Tim has checked in more tests. We are going to use the tests written for hermes and cerberus read accesses in the 'How to debug Euterpe' presentation on monday. We will be looking for the time it takes from the instruction issue, entry to nb, and completion of access.
 05/24 Tim is going to send lisar a list of the test names that he intends to write so they can be incorporated into the test template.

Suspended Items

Item: Unsnap code
 Who: lisar, jeffm
 Status: Suspended

02/15 The issue of restarting the hardware from an IKOS dump was discussed and the need for an architectural snap/unsnap facility was questioned.
 Since the meeting it has been re-discovered (jeffm wasn't there to remind us of an earlier decision) that we are planning on loading architectural state into an IKOS simulation and not from a total IKOS logic dump.
 We also determined that when it came time to run some of the larger tests (real-time benchmark) we would need the capability to start an IKOS simulation from an architectural dump anyhow.
 03/01 For the short term we are going to focus on a simpler approach for loading and running DVTs, the kernel, and kernel tests. This item will likely come back in April.
 05/10 Back to life. Does the IKOS support RAM dump?
 05/24 This item was suspended again. There are no resources allocated to this item at this time.

Item: When do we have a full calliope simulation available (IKOS)?

Who: lisar

Status: Suspended.

04/26 This topic was raised as we talked about when the Snap/Unsnap item should be brought back from the Suspended list.

05/10 Lisar was not able to attend the meeting.

05/24 Suspended. This is related to the snap/unsnap item (above).

Completed Items

Item: Preparation for 'How to Debug Euterpe' presentation

Who: jeffm, doi, gregg, others?

Status: Done.

05/10 Need to put together a packet of information about the tools and utilities used for effective debugging.

std tool (create source listing annotated with trace information)

mkimg (building load images and files suitable for loading in the HW simulators)

likelevel log description

... and more.

05/17 Jeffm and Lisar did a great job of this presentation.

There is an html document to server as an aid. The document can be found in /u/chip/euterpe/doc/debug/EuterpeDebug.html.

Item: Modify tests in diag tree to use tcc instead of tgcc

Who: guarino, jeffm, doi

Status: Done.

03/29 Loretta has changed the diag tests to use tcc instead of tgcc but has not checked in the changes.

Lisar wants these changes to be made after we are able to run the tests successfully using tgcc.

04/20 Loretta has the stand/diag tests converted to use TCC.

Derek is going to modify the Makefiles in stand/diag/memhi so that the programs that required tgcc will have an explicit rule.

04/26 Loretta is ready to check in the TCC related changes. Derek is to check with lisar if this is OK.

05/03 Loretta was ready to check in the changes but a last minute test uncovered a compiler bug. Checkin will happen after the tests build and run OK.

Terp Feature Status

done o Add support for host I/O through the sdram

inprog o Ifetch protection granularity
- performance vrs accuracy tradeoff

inprog o Fetch instructions as octlets

inprog o Accuracy wrt HW simulator(s?)

o Better latency model for Calliope accesses

o Implement hardware configuration through Cerberus regs (SDRAM parameters, dram enable?)

o Checkpoints/Snapshots

done o hermes timeout machine checks

- does jeff have a test for this? Currently the TSA is followed (immediate mchk) instead of waiting for watchdog timeout.

inprog o ability for terp to load hermes sections

- lisar would like this functionality added

inprog o remove stbio from hwterp.

Performance Test Status

Here is the current (un-prioritized) list of desired performance analysis.

Tests to be created:

done 1) stores/loads to dram, hermes, cerberus, load from rom
done 2) icachemiss
done 3) dcachemiss
4) load ltlb entry (write+read)
5) load gtlb entry (write+read)
6) NB overflow
7) generate an interrupt
8) return from interrupt
9) multiple cylinders trying to take exceptions at the same time
10) predicated branch
11) unpredicted branch
inprog 12) One instruction from each instruction class:

arith	store	eshifty	imul64	
branch		swap	epop	idiv
branchx		stored	imul4	gfmul
load	eset		imul8	bgate
loadx	eshift		imul16	atomic ops
nbload		eshiftx	imul32	

done 13) Inner loop sequences:
cable_in_main_handler-- inner loop in EQ_UPDATE_WEIGHTS() (khp)
IDCT code (fur)
pseudo-Huffman decode loop (fur)
a reconstruction routine from macroblock.c (fur)
NTSC encode loop (fur)
rs_compute_syndrome() (fur)

Test Status and General Discussion

The nullTest passed (last week, actually)!

Loretta is going to send doi a list of the ukernel tests that could be run next. Doi will then make sure that we can build preloaded versions of these tests (like the nullTest).

The ggmul8 will likely have a constraint that the src and dst registers can not be the same. The compiler is going to be modified to prevent this condition from being emitted and the assembler is going to be modified to emit a message if this does in fact occur.
----- End of forwarded message -----

From: lisar (Lisa Robinson)
Sent: Thursday, May 25, 1995 12:50 PM
To: claseman; jeffm
Cc: guarino; gmo; doi; gregg; lisa; tbr; mws; billz; dickson; woody
Subject: Re: Software Bringup Meeting Minutes - May 24, 1995

First of all I should apologise for not making the meeting. It completly slipped my mind.

As far as the performance cases, I don't think that this should be on hold I have run 4 cases out of the 6 through and they all show a discrepancy with respect to the number of cycles it takes on terp. I would like for jeffm (or a designate in the design group) to say this is what is expected of the HW and have the tests fail (on either terp or the HW) when it isn't in that range.

I am not putting this as the highest priority for running against the HW as I believe that much of the work can be done by looking at the test and saying that (from HW perspective) this is the number of cycles it should take and then running it on terp. I do know that terp is optimistic and that is what worries me.

Lisa R.

	terp	euterpe	
hermes_perf	0x64	0xba	
cerb_perf	0x50		This must be optimistic!
dram_perf	0xaa	0xce	
rom_perf	0x82	0x2f4	

Now of course any result has to be scaled to take into account that on the HW I'm running cerberus only 5 time slower than the sofa.

In addition the hermes device responds differently from calliope (ie buffer/ registers).

I think that there is alot to do here.

Lisa R.

From: hopper (Mark Hofmann)
Sent: Thursday, May 25, 1995 5:08 AM
To: tbr (Tim B. Robinson)
Subject: cpu cycles for CAD tapeout

Hi Tim,

The CAD folks have been pushing around the idea of adding hardware to speed up back end tapeout processing. The conclusion is that breaking a chip into strips is probably more difficult than it seems at first blush and doesn't address memory problems. What does seem promising is to make Vlsimm run in parallel. If tom added this (it is something he's thinking about) and we had a multi-processor machine with shared memory we might see a significant gain.

Therefore, we thought of commissioning the sysadm group with two tasks:
1. find out specs of the Power Challenge: how many processors, how much memory, how much swap, how much memory per process, is anything 64-bit, how does it compare with our current challenge? 2. same questions of the DEC alpha.
Both of these machines would speed up DRC, waffle, synthesis and fracture.
In addition, a true 64-bit architecture might speed up LVS (which is not a bottleneck, at present).

Comments?

-hopper

From: doi (Derek Iverson)
Sent: Thursday, May 25, 1995 12:03 PM
To: guarino; gmo; doi; gregg; claseman; lisa; jeffm
Cc: hestia
Subject: Software Bringup Meeting Minutes - May 24, 1995

Software Bringup Meeting

May 24, 1995

Next Meeting: May 31 at 2:00 pm.

Attendees: guarino, gmo, doi, gregg, claseman, lisa, jeffm

New Action Items

None.

Review of Action Items

Item: Modify startup code in stb/stand to read the cerberus node number

Who: gmo

Status: Pending

05/03 No new progress

05/10 Ditto.

05/24 Ditto.

Item: Plan for testing remote debugging environment

Who: everyone

Status: Pending

04/13 This will be discussed at the next meeting.

04/20 Short discussion on the pieces required and if we want to have a standalone version of the hostio software. Wally has started on this already.

04/26 Any work with snoopy/dram models will begin after euterpe tapeout when jeff has more time.

05/03 Wayne F. expects to have a prototype ISA card ready for testing tomorrow. Gmo is in the process of writing code to test the card.

05/10 Gmo has written a test program. Gmo and Wayne need to get together and try it on the board.

05/24 Wally has added more connectivity between the ukernel and the debugger (ability to interrupt kernel via the cerberus 'forced interrupt' bit).

Gmo is ready with a test when Wayne has the board ready.

Item: ukernel needs to detect machine checks and 'do the right thing'

Who: lisa

Status: Pending

04/13 No new progress.

04/20 On list of things to do. Lower priority.

04/26 Ditto.

05/03 Ditto.
05/10 Ditto.

Item: Create performance test plan

Who: claseman

Status: [11/30] No progress as focus is on functionality.

- 11/30 We continue to run tests to help us compare terp vs hardware performance.
We still need to put together the actual performance tests that need to be run on the hardware.
- 05/10 Back to life. Tim has checked in more tests. We are going to use the tests written for hermes and cerberus read accesses in the 'How to debug Euterpe' presentation on monday. We will be looking for the time it takes from the instruction issue, entry to nb, and completion of access.
- 05/24 Tim is going to send lisar a list of the test names that he intends to write so they can be incorporated into the test template.

Suspended Items

Item: Unsnap code

Who: lisar, jeffm

Status: Suspended

- 02/15 The issue of restarting the hardware from an IKOS dump was discussed and the need for an architectural snap/unsnap facility was questioned.
Since the meeting it has been re-discovered (jeffm wasn't there to remind us of an earlier decision) that we are planning on loading architectural state into an IKOS simulation and not from a total IKOS logic dump.
We also determined that when it came time to run some of the larger tests (real-time benchmark) we would need the capability to start an IKOS simulation from an architectural dump anyhow.
- 03/01 For the short term we are going to focus on a simpler approach for loading and running DVTs, the kernel, and kernel tests. This item will likely come back in April.
- 05/10 Back to life. Does the IKOS support RAM dump?
- 05/24 This item was suspended again. There are no resources allocated to this item at this time.

Item: When do we have a full calliope simulation available (IKOS)?

Who: lisar

Status: Suspended.

- 04/26 This topic was raised as we talked about when the Snap/Unsnap item should be brought back from the Suspended list.
- 05/10 Lisar was not able to attend the meeting.
- 05/24 Suspended. This is related to the snap/unsnap item (above).

Completed Items

Item: Preparation for 'How to Debug Euterpe' presentation

Who: jeffm, doi, gregg, others?

Status: Done.

- 05/10 Need to put together a packet of information about the tools and utilities used for effective debugging.

std tool (create source listing annotated with trace
information)
mkimg (building load images and files suitable for
loading in the HW simulators)
likelevel log description
... and more.

05/17 Jeffm and Lisar did a great job of this presentation.
There is an html document to server as an aid. The document
can be found in /u/chip/euterpe/doc/debug/EuterpeDebug.html.

Item: Modify tests in diag tree to use tcc instead of tgcc

Who: guarino, jeffm, doi

Status: Done.

03/29 Loretta has changed the diag tests to use tcc instead
of tgcc but has not checked in the changes.
Lisar wants these changes to be made after we are able to
run the tests successfully using tgcc.

04/20 Loretta has the stand/diag tests converted to use TCC.
Derek is going to modify the Makefiles in stand/diag/memhi
so that the programs that required tgcc will have an
explicit rule.

04/26 Loretta is ready to check in the TCC related changes. Derek is
to check with lisar if this is OK.

05/03 Loretta was ready to check in the changes but a last minute test
uncovered a compiler bug. Checkin will happen after the tests
build and run OK.

Terp Feature Status

```

done  o Add support for host I/O through the sdram
inprog o Ifetch protection granularity
      - performance vrs accuracy tradeoff
inprog o Fetch instructions as octlets
inprog o Accuracy wrt HW simulator(s?)
      o Better latency model for Calliope accesses
      o Implement hardware configuration through Cerberus regs
        (SDRAM parameters, dram enable?)
      o Checkpoints/Snapshots
done  o hermes timeout machine checks
      - does jeff have a test for this? Currently the TSA is
        followed (immediate mchk) instead of waiting for watchdog
        timeout.
inprog o ability for terp to load hermes sections
      - lisar would like this functionality added
inprog o remove stbio from hwterp.
  
```

Performance Test Status

Here is the current (un-prioritized) list of desired performance analysis.

Tests to be created:

```

done  1) stores/loads to dram, hermes, cerberus, load from rom
done  2) icachemiss
done  3) dcachemiss
      4) load ltlb entry (write+read)
      5) load gtlb entry (write+read)
      6) NB overflow
      7) generate an interrupt
      8) return from interrupt
      9) multiple cylinders trying to take exceptions at the same time
      10) predicated branch
  
```

11) unpredicted branch
inprog 12) One instruction from each instruction class:

```
arith    store    eshift    imul64
branch   swap     epop      idiv
branchx  stored    imul4     gfmul
load     eset     imul8      bgate
loadx    eshift   imul16     atomic ops
nbload   eshiftx  imul32
```

13) Inner loop sequences:
done cable_in_main_handler-- inner loop in EQ_UPDATE_WEIGHTS() (khp)
IDCT code (fur)
pseudo-Huffman decode loop (fur)
a reconstruction routine from macroblock.c (fur)
NTSC encode loop (fur)
rs_compute_syndrome() (fur)

Test Status and General Discussion

The nullTest passed (last week, actually)!

Loretta is going to send doi a list of the ukernel tests that could be run next. Doi will then make sure that we can build preloaded versions of these tests (like the nullTest).

The ggfmul8 will likely have a constraint that the src and dst registers can not be the same. The compiler is going to be modified to prevent this condition from being emitted and the assembler is going to be modified to emit a message if this does in fact occur.

From: vanthof (vant)
Sent: Thursday, May 25, 1995 8:51 AM
To: tbr (Tim B. Robinson); hopper (Mark Hofmann); geert (Geert Rosseel); vo (Tom Vo); lisar (Lisa Robinson)
Cc: vanthof (Dave Van't Hof); wampler (Kurt Wampler); tom (Tom Laidig)
Subject: lvs results

Well, the lvs finished and it's still not very clean:

NUMBER OF UN-MATCHED SCHEMATICS DEVICES	=	5245
NUMBER OF UN-MATCHED LAYOUT DEVICES	=	5246
NUMBER OF MATCHED SCHEMATICS DEVICES	=	2091442
NUMBER OF MATCHED LAYOUT DEVICES	=	2091442

The device counts are what I expected, but there still seems to be some problems. Of course, I complicated the issue by munging the node names to something completely unreadable (I do have a translation table though...) which will slow down the debug.

I'll let you know what I find.

The output is in:

/u/vanthof/compass/mobi/euterpe/tapeout/euterpe.compare/euterpe.lvs

The nodename translation file is:

/u/vanthof/compass/mobi/euterpe/tapeout/euterpe.compare/short.names

Thanks,
Dave

--
Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100
"Don't blame me! I didn't vote for him"

From: tbr
Sent: Tuesday, May 23, 1995 9:39 PM
To: stick (Bruce Bateman)
Cc: billz; bpw; dickson; geert; mws; wingard; woody
Subject: Re: Custom Cronus Blocks - Cache/Tag

Bruce Bateman wrote (on Tue May 23):

>
> Yes! It's really important we do not let such an optimisation force
> us to make changes in the sofa (other than the realtively easy step of
> moving the flops from the sofa into the array). If the extra delay in
> the sofa costs us an extra cycle, we will not only have created a
> major schedule hiccough from the necessary redesign of the pipeline,
> but we may well have lost more system level performance than we would
> from simply backing the cycle time down a little to accommodate what
> you can deliver.

My concern was being able to make the cache meet a two cycle operation.
I did not consider slowing the clock, 'cause the marching orders were
400MHZ.

It's appropriate for the large arrays to be a pacing item, but we have to be able to match
it with other critical paths to be able to use the performance

> It's not clear to me why even with two decodes we
> cannot have them at essentially the same physical location (ie in the
> center).

This is not practical from a layout perspective, unless you want to give
up area. Putting two read-decoders in the center would force the use of
two rather than one write decode - one on each side (there's no way we
could fit both the read AND write decoders in the middle). Then we would
have 4 discrete decoders rather than the current three.

OK, that clarifies.

>
> Can we please make sure that all interested parties get advised when a
> change this major is contemplated, so that we do not get a surprise
> after the fact?
>

I did discuss this early on with Drew. I didn't realize that it would be
a major problem for the sofa. My appologies. What now? Do we want to
change back to read in the middle and write on the sides?

Since the floorplan is still fluid it's not clear it is major problem, it's just one thing
we have to be aware of and design around. On Euterpe, it's definitely one of the most
critical paths and we had to work hard to meet it, which would not have been possible with
the extra length of having the decoder at the outside. Of course had it been there we'd
have been re-arranging the floorplan.

So don't change anything unless we get to the point where we can't find a viable
floorplan.

Tim

From: stick (Bruce Bateman)
Sent: Tuesday, May 23, 1995 11:19 AM
To: tbr
Cc: billz; bpw; dickson; geert; mws; wingard; woody
Subject: Re: Custom Cronus Blocks

> Date: Mon, 22 May 1995 22:20:06 -0700
> From: tbr (Tim B. Robinson)
> To: stick (Bruce Bateman)
> Cc: billz, bpw, dickson, geert, mws, wingard, woody
> Subject: Re: Custom Cronus Blocks
>
> Bruce Bateman wrote (on Mon May 22):
>
> In the read-port, the data being addressed goes into a clocked
> sense-amp (pre-charged high) on every rising clock edge, which then
> feeds an SR-latch. The SR-latch holds the previous data valid during
> the pre-charge state of the sense amp and thus acts like a poor man's
> holding register. The difference is that new data appears on the
> output after the first rising clock after the cominatorial delay
> through the cache/tag array. Thus, at slow clock speeds, the new
> data may appear after one clock cycle instead of two clock cycles
> at higher clock speeds.
>
> This is *bad* because it means the logical behavior changes with clock
> frequency and the surrounding logic cannot rely on your "holding"
> register to be holding anything. In the current Euterpe (and hence
> Cronus) design, for example on the dcache output, we have a two cycle
> path from the dcache, through NB and into the XLU. I'd expect this to
> be a critical path in the cronus layout also, so we will need
> something stable for two clocks. Obviously adding a true holding
> register on the outside will add additional delay.
>
> ...<snip>...
>
> The pinouts are essentially the same as the euterpe cache/tag except the
> physical position of the read and write port decoders has been swapped,
> so that there are two sets of read-port address pins and only one set of
> write port pins. Thus the pins on the cache are:
>
> Is this essential? It will mean that the critical read path is
> penalized instead of the non-critical write path, since the read
> address will have significantly longer wire load than with both copies
> at essentially the same location.
>

The current design meets a 5ns (2 clock cycles at 400MHz) cycle with only 2-300ps set-up to clock in the sense amp. In just the last week I've managed to improve this to the current number. Before I only had 1-200ps margin. I can add some sort of latching mechanism to the output at the expense of clock to output time, which is currently running slightly over 900ps into 100ffd. Changing to a central read port decode will cost me 2-300ps of additional delay thru the cache and will essentially erase all my margin in set-up time to the sense amp.

BB

From: hopper (Mark Hofmann)
Sent: Tuesday, May 23, 1995 1:24 AM
To: Derek Iverson
Cc: ken (Ken Hsieh); sysadm
Subject: Re: Request for help with installing a Verilog patch

Derek Iverson writes:

Hopper has some new verilog sources that I would like to install in /a/cadence_rel/rel.9404.sun4.

Here is what I would like to do....

```
cd /a/cadence_rel/rel.9404.sun4
mv data data.before.s016

cd tools.sun4/verilog
for i in README etc examples include lib src ; do
    mv $i $i.before.s016
done

cd ../..
zcat ~hopper/vendor/cadence/verilog/verilogxl02.10-s016sun4.t.Z | tar -xvf -
```

It turns out that of the entire list of crap they stuff on a tape, the following is a list of all the files that have actually changed.

Because the list is so short, we may just want to extract these specific files (making a copy of the original first, of course) instead.

Comments?

```
./data/defaults
./data/desc
./data/submittor
./tools.sun4/verilog/lib/libcosim.a
./tools.sun4/verilog/lib/libvoids.a
./tools.sun4/verilog/lib/libstdfa.a
./tools.sun4/verilog/lib/vlog.o
./tools.sun4/verilog/lib/libdld.a
./tools.sun4/verilog/include/acc_user.h
```

Yes, let's just extract the deltas between the 2 versions. We can wait for a full release-9502 or whatever- before doing a wholesale copy.

After this is done, I will want to do something similar to the HP release. It turns out that we only have pointers to the 9403 release of verilog for the HP machines in /a/cadence_rel.

Ken,
Do you know if we have the 9404 software for the HP machines?

Thanks,
doi

-hopper

From: tbr
Sent: Tuesday, May 23, 1995 12:24 AM
To: doi (Derek Iverson)
Cc: hopper; ken; sysadm
Subject: Request for help with installing a Verilog patch

Derek Iverson wrote (on Mon May 22):

Hopper has some new verilog sources that I would like to install in /a/cadence_rel/rel.9404.sun4.

Here is what I would like to do....

```
cd /a/cadence_rel/rel.9404.sun4
mv data data.before.s016

cd tools.sun4/verilog
for i in README etc examples include lib src ; do
    mv $i $i.before.s016
done

cd ../..
zcat -hopper/vendor/cadence/verilog/verilogx102.10-s016sun4.t.Z | tar -xvf -
```

It turns out that of the entire list of crap they stuff on a tape, the following is a list of all the files that have actually changed.

Because the list is so short, we may just want to extract these specific files (making a copy of the original first, of course) instead.

Comments?

```
./data/defaults
./data/desc
./data/submittor
./tools.sun4/verilog/lib/libcosim.a
./tools.sun4/verilog/lib/libvoids.a
./tools.sun4/verilog/lib/libsdfa.a
./tools.sun4/verilog/lib/vlog.o
./tools.sun4/verilog/lib/libdld.a
./tools.sun4/verilog/include/acc_user.h
```

After this is done, I will want to do something similar to the HP release. It turns out that we only have pointers to the 9403 release of verilog for the HP machines in /a/cadence_rel.

Ken,

Do you know if we have the 9404 software for the HP machines?

Do we know who's still running verilog on the HP's? I doubt they run it faster than the sparc 10's so are we spending time maintaining something we no-longer need?

Tim

From: stick (Bruce Bateman)
Sent: Monday, May 22, 1995 9:40 PM
To: bpw; billz; mws
Cc: tbr; geert; dickson; woody; wingard
Subject: Re: Custom Cronus Blocks - Cache/Tag

> Date: Mon, 22 May 1995 19:23:26 -0700
 > From: mws (Mark Semmelmeier)
 > To: bpw, billz, stick
 > Subject: Re: Custom Cronus Blocks - Cache/Tag
 > Cc: tbr, geert, mws, dickson, woody, wingard
 >
 > Thanks for the detailed info, Bruce.
 >
 > > From stick Mon May 22 17:28:45 1995
 > > The
 > > address inputs use a mux-flop holding reg, with the mux select pin
 > > used to "hold" the address valid when active low. For a read (load)
 > > cycle, the address must be held for two cycles. For a write (store)
 > > cycle, the address must be held for four cycles.
 >
 > I assume you mean the cache will hold them for us based on the outside
 > logic driving the hold control correctly.

That is correct.

> > In the read-port, the new
 > > data may appear after one clock cycle instead of two clock cycles at
 > > higher clock speeds. The read port timing will thus look something
 > > like:
 >
 > > phi_am _____/_____/_____/_____/_____/_____/_____/_____
 > > rhld_bnm _____/_____/_____/_____/_____/_____/_____/_____
 > > radd_bm XXXXX_____A1_____XXXXXXXXXXXXXXX_____A2_____XXXXXXXXXXXXXXX_____
 > > out_bm XXXXXXXXXX_____D0_____XX_D0_or_D1_XXX_____D1_____XX_D1_or_D2_XX_____
 >
 > The 1-tick valid output is not a problem for the I Cache or I Tag,
 > which didn't use HR output registers. The D Tag can easily be fixed
 > to replace its current HR emulator. The D Cache is more of a problem
 > because its output register was an HR to drive 2 ticks through a mux
 > (not muxflop) in NB and then all the way into XLU. We will either
 > have to convert the data out pipe from half rate to full rate with a
 > flop resting point in NB or the cache will have to hold its output for
 > 2 ticks. Drew mentioned that he might be interested in having you
 > gate the output RS register to save power since it really only needs
 > to clock half the time. If that is done, then we are back to 2-tick
 > stability for DCache, and the other usages will accept that too.

We need to discuss this ASAP is we want to change the current design.


> We also have to be careful about drive strength. With the output
 > registers in the array, we lose SOFA adjustment of drive strength.
 > I hear that the cache is designed for 100ff of load. This will almost
 > certainly require an extra rank of dout buffer for DCache and ITag in
 > a critical timing path.
 > For ICache, we can probably use the bypass mux to serve that purpose.
 > DTag dout is not timing critical.
 >

>
> > The write port addresses are synchronous, as mentioned above, but the
> > data-in and write-enables are async, same as in the euterpe cache/tag.
> > Thus the write port timing looks something like:

> As long as we have so much synchronous already, is it better to make
> it all synchronous? Are we going to be able to control WE timing
> easily enough in Synopsys+AutoPlace+AutoRoute methodology?
>

I asked you about this several weeks ago and you input was since the euterpe cache already
tollerated the async inputs, why complicate the din/we inputs on the cache unnecessarily.
Thats why I left them this way.

> It seems we are not exploiting the address hold facility in the
> picture below. Is the facility still usable to hold?



Yes, my error in the diagram.

> > The pinouts are essentially the same as the euterpe cache/tag except
> > the physical postition of the read and write port decoders has been
> > swapped, so that there are two sets of read-port address pins and
> > only one set of write port pins.
>

> Two read ports on opposite sides of a big structure is bad news for
> our critical read address path. Can we get back to one read address
> target to cut our path length and loading? Drew also suggests that
> the port cycling twice as fast (read) should be done in a more
> centralized location to save power.
>

I wouldn't want to change at this point. The reason I use two decodes on the read port
was to save me a fanout of 2 in the critical read path (i.e.
each read-decoder drives 1/2 wordline). This was necessary becuae the cache is very
tight for meeting the 2-cycle latency. Of course, this pushed the problem of the extra
loading onto you in the sofa.

BB

From: doi (Derek Iverson)
Sent: Monday, May 22, 1995 8:16 PM
To: ken
Cc: sysadm; hopper
Subject: Request for help with installing a Verilog patch

Hopper has some new verilog sources that I would like to install in
/a/cadence_rel/rel.9404.sun4.

Here is what I would like to do....

```
cd /a/cadence_rel/rel.9404.sun4
mv data data.before.s016

cd tools.sun4/verilog
for i in README etc examples include lib src ; do
    mv $i $i.before.s016
done

cd ../..
zcat ~hopper/vendor/cadence/verilog/verilogx102.10-s016sun4.t.Z | tar -xvf -
```

It turns out that of the entire list of crap they stuff on a tape, the following is a list
of all the files that have actually changed.

Because the list is so short, we may just want to extract these specific files (making a
copy of the original first, of course) instead.

Comments?

```
./data/defaults
./data/desc
./data/submittor
./tools.sun4/verilog/lib/libcosim.a
./tools.sun4/verilog/lib/libvoids.a
./tools.sun4/verilog/lib/libsdffa.a
./tools.sun4/verilog/lib/vlog.o
./tools.sun4/verilog/lib/libdld.a
./tools.sun4/verilog/include/acc_user.h
```

After this is done, I will want to do something similar to the HP release. It turns out
that we only have pointers to the 9403 release of verilog for the HP machines in
/a/cadence_rel.

Ken,
Do you know if we have the 9404 software for the HP machines?

Thanks,
doi

From: paulp (Paul Poenisch)
Sent: Monday, May 22, 1995 1:05 PM
To: Mark Hofmann
Cc: cadettes; geert (Geert Rosseel); tbr (Tim B. Robinson)
Subject: Re: Pads

hopper writes:

>
> vant writes:
> Geert,
> The types of rules that Paul is referring too are 'probabalistic
> failure mechanisms'. Which means that in certain cases, it might fail but
> then, it might not.
>
> I agree with you that unless it is checkable, then we can't ensure the
> chip to work.
>
> Paul is still working on a list of things which are recommendations and
> putting them in a 'style guide'. I only know of one rule right now that is
> uncheckable, and that is 'fat metal spacing'.
> This rule is complex and causing
> the metal synthesis to increase drc run times by about 3x, and then backend
> tapeout to increase by many days.
>
> I will be out most of the weekend, so I don't know how much I can help,
> but I'll try.
>
> Thanks,
> Dave
>
> Geert,
>
> To amplify this: It has been agreed that there are to be 2 documents:
> 1. DRC Bible 2. Style Guide. All circuits must pass all rules in the
> the DRC Bible. For a rule to be placed in the DRC Bible it must be checkable.
> There are many steps which Al and others can not give hard and fast
> rules for, but which they feel if done a certain way, will increase yield.
> These things are to be placed in a separate Style Guide. Paul is
> working on this. The CAD group is going to first focus on getting the
> DRC Bible fully and correctly implemented. We will work on the style
> guide second. (A Style Guide example: Uniform spacing of certain vias. How do you define
> "uniform"?
> When is there a violation?. We will need to implement certain area
> density functions to check this.)
>
> It is my understanding- and Paul please correct me if I err- that
> all circuits which pass the full DRC Bible ruleset should yield
> functioning die, to the best of the processing group's knowledge,
> without the need for any further "style" or "visual" checks.
>
> -hopper
>

All,

It is my intention that when the design rules have fully stabilized any design that meets those rules will yield at a reasonable level. Given that as a basic underlying goal there are several caveats. First, the design rules are not yet stable, there are no doubt rules that are or will be needed for the final process flow that we freeze that are not in the rules now. Additionally there are probably rules that we don't or will not need that are now being followed.

Second I don't know right now what "yield at a reasonable level" actually means, it could

be 75% or 25%, we are currently targeting for 50% but there are too many variables to predict it right now. Third I'm not sure when the design rules will be fully stabilized, maybe by the end of the year but it could be the middle of next year before things fully settle down (remember that this is a completely new process running in a new fab, there are more variables than you can shake a stick at).

As for the style rules, I think everyone will agree that in any process there are designs that yield well and those that don't. Often no one ever finds out exactly why a design yields poorly, production of the design is just stopped because the competition has killed it. The designs that yield poorly will almost always yield as well as the "good" designs when the fab is having a good day, on a bad day the good design still yields OK but the poor design zeros.

Having worked with the design, layout and CAD personnel at this company for the last two years or so I believe that we have a very good group that routinely will produce "good" designs given that they have an understanding of what a good design is. The intent of the style rules is to give the people who need them a set of guide lines as to what constitutes good design practice. This is particularly important in that the process we are running here is so different from those used elsewhere, what is good practice somewhere else is a disaster here.

Finally, the current set of design rules (5.0, which you don't have yet) has known holes in it. There are structures in the current I/O cell which will pass these design rules but will cause almost zero yield (the clamp transistor for instance). Currently I do not know how to write a rule that will prevent structures like these from being designed without gutting the process and because I can't write the rule the CAD group can't write a DRC for it. I expect that I will need to talk to the design and/or layout people that will be redesigning the I/O cells about how to avoid these problems (style) but I believe that in the next couple of months we will be able to write the needed rules and they can be incorporated in future designs. Hopefully in the meantime style will be enough, it's all we can do for now.

Paul.

From: William Herndon [bill@polyhymnia.microunity.com]
Sent: Monday, May 22, 1995 12:05 PM
To: graham@polyhymnia.microunity.com
Cc: wingard@polyhymnia.microunity.com; dane@polyhymnia.microunity.com;
 graham@polyhymnia.microunity.com; geert@polyhymnia.microunity.com
Subject: Re: Summary of Cronus/Atlas Meeting on May 19

We also need clock driver and flavors layout help. I am concerned about jumping in iobyte without a coherent plan. Dane has done the vrr and process code circuits. We also have a clock delay generator for the quadrature, all the input latches, and getting from ecl to cmos levels. The delay generator and ecl latches have a lot of work done on them. No work on the ecl to cmos level translator has been done, no work on how to clock the fifo etc. My point is that there is a lot to do on the iobyte. I will focus my efforts on getting the clock driver circuit flavors ready for layout.

On Mon, 22 May 1995 graham@polyhymnia.microunity.com wrote:

```

> Drew, I see from the summary that I/O byte is behind schedule.
> Dane tells me that his schematics are completed, and that the issue is
> layout resource. Should we locate a layout contractor - or have the
> circuit designers do some layout?
>
> Please let me know how we can assist to regain the schedule.
> Regards, Graham.
>
> > From: wingard Sun May 21 23:06:45 1995
> > Date: Sun, 21 May 1995 23:06:32 -0700
> > From: wingard (Drew Wingard)
> > To: cronus
> > Subject: Summary of Cronus/Atlas Meeting on May 19
> > Content-Length: 8502
> >
> > Hi,
> >
> > Here is a summary of last weeks Cronus/Atlas Status meeting
> >
> > Attendees: tbr, lisar, trancy, pulp, bill, stick, vo, ong, hopper, wampler,
> > brianl, fwo, wingard
> >
> >
> > Note: Lines beginning with ">" are from Geert's message (of 5/9) reporting the
> > results of the meeting on May 5th.
> >
> > > 1. Baseplate
> > >
> > > Current status : We have a gards-compiled cronus baseplate
> > > which can be used for routing experiments,
> > >
> > > Plan : By May 31, we want to have a final baseplate.
> > >
> > > Action items : * Decide on final die size : Drew
> > > STATUS - Makefile and tools nearly in place to support this
> > > activity. Should make June 1 deadline, but in
> > > some danger of slipping.
> > >
> > > * Pading assignment : Warren
> > > STATUS - On hold pending floorplanning.
> > >
> > > * Floorplan, custom block
> > > placements : Warren
> > > STATUS - On hold pending die size work. Power interface
  
```

> > (hemming) cell work is ongoing. First version of
> > clock mast cut-outs released.

> > > * Clock Tree generation : Kurt, Bill
> > STATUS - In progress. Bill has communicated plans for a
> > single-mast clock system to Kurt. Kurt indicates
> > on-schedule completion is very likely.

> > > * Build dummy cells for
> > TTL I/O and : B.P
> > IOBYTE : Dane
> > STATUS - No progress to report.

> > > 2. Custom blocks

> > > Current status : GTLB done
> > Register file : end of this week
> > Caches : layout at top-level
> > Tags : layout has not started yet
> > NB : layout has just started
> > TTL I/O : not designed yet
> > IOBYTE : not started yet

> > > Plan : finish all custom blocks by end of May
> > except TTL I/O and IOBYTE

> > > Action Items : * Register File : B.P, Mike, Orlando
> > STATUS - Polygons finished on 5/19. Largely DRC/LVS clean.
> > Top-level should be LVS clean by 5/24.

> > > * Caches : Bruce, Efelias, Dtacmo
> > STATUS - Laying out 'last' two cells. Final size is known.

> > > * Tags : Bruce, Efelias, Dtacmo
> > STATUS - No progress, but can use only the cache leaf
> > cells and therefore be only top-level assembly.

> > > * NB : Vikki, B.P.
> > STATUS - Polygons should be finished by about 5/24.

> > > * TTL I/O : B.P., ??
> > STATUS - Design goals specified. In design. Some early
> > layout work has begun?

> > > * IOBYTE layout can start by end of next week.
> > We'll need schematics by then : Dane, Bill
> > STATUS - No progress was reported.

> > > New Action Items : * IOBYTE is apparently slipping badly. It will not
> > make the June 1 deadline. Drew to find out what we
> > can do to make better progress.

> > > * XLU: assumption was that we would try to build XLU
> > datapath out of Atlas std cells (with 'custom'
> > placement and hand route). Tbr to ask Karzes about
> > modifying generation program to output std cells.
> > Drew to consider adding mux16 to support Stage 2.

> > > * CERPOKGEN: Power OK circuit needs to be designed.
> > Since the caches will be a client, consider using
> > circuit that Bruce had started designing.

> > > 3. Atlas database

> > > Current status : Database builds from toplevel. An initial version
> > of about everything is there.
> > STATUS - Atlas database builds from scratch as of 5/18

Plan : Have a complete and accurate Atlas database by the end of May
In building Cronus sub-blocks, everybody should point to
/u/chip/atlas and not local versions. That will catch
database problems early.

Action Items : * Review XL, XS cells
timing, layout, cap models : Warren
STATUS - All cells have passed corner analysis. Some
had to be resized for better margin. Latches
and muxes have been performance tuned. Other
cells ongoing.

* Finish timing numbers : Fred
STATUS - No progress to report. Hopper to help with
capacitance and input loading work.

* Complete verilog libraries : Brianl
STATUS - Basic libraries installed. Still need models
for Cerberus and 'SC' cells.

* Complete Synopsis libraries : Brianl
STATUS - Basic libraries installed. Severe performance
degradation with latest version of Synopsys
noticed, but can still use prior version until
fix/workaround is provided.

New Action Items : * Enhance Verilog libraries : Brianl and Drew
Tbr has provided a list of 'missing' cells
that prevent him from compiling Cronus for
simulation at the top level. We need to get
the truly needed ones installed ASAP.

* Enhance Gate Family : Warren and Drew
We need XOR and MAJ gates, and perhaps an
asynchronously-loadable latch for Cerberus.

> 4. Cronus/verilog/bsrc, Makefile.rules, GARDS

Current status : Initial versions of Makefile.rules exist. We can
build a sub-block using the atlas database on
the checked in baseplate model.

Plan : Have a "good" Makefile.rules that works well enough that
other people can start mapping Euterpe blocks by the end of
May
Implement " put the top-level together" strategy in the
equivalent of euterpe/verilog/bsrc/Makefile.tst

Action Items : * Makefile.rules : Brianl
STATUS - "Better" version released. Ongoing.

* GARDS model additions to deal
with power and clock obs : Tau, Kurt
STATUS - Completed and tested - works great!

* Pim2pif upgrade : Hopper
STATUS - Completed.

* Take euterpe snapshot : Tbr
STATUS - Completed. Cronus.V released.
Tbr has been cutting out Euterpe blocks that
Cronus will not have, and bringing over
behavioral simulation models of the custom
blocks. Billz has swapped in a behavioral
NB buffer description. Dickson has been

> > modifying the Cerberus register map to remove
> > bits associated with unused functionality.

> > > * Floorplanning tools : Drew
> > STATUS - Makefile and tools nearly in place to support this
> > activity. Should make June 1 deadline, but in
> > some danger of slipping.

> > > * verilog/bsrc/Makefile : Drew
> > STATUS - See above. Rest of this activity is largely in
> > the hands of Tbr and Brianl.

> > > 5. "Implementation of mapping"

> > > No plans for this month
> > STATUS - Brianl has been 'test' mapping the top-level
> > Euterpe blocks to identify mapping and P&R issues.
> > This should also provide valuable feedback in
> > setting the die size.

> > > 6. Packaging

> > > Current status : We have "a plan" (read all about it in Mosaic)

> > > Plan : By the end of May, make decisions on all aspects of "the plan"

> > > Action Items : * Call a set of meetings
> > > during May to turn our
> > > plan into a set of decisions : Geert
> > STATUS - Met on 5/15 and decided on 70mm TAB frame.
> > Trancy investigating having someone else do
> > ILB/OLB work, but lead times are about as
> > bad as doing it ourselves. Drew spoke with
> > MMS and got clarification on module spacing
> > requirements. Trancy needs firm substrate size
> > to proceed much further. Drew to provide that
> > this week.

> > > 7. Test

> > > Current Status : Mudge & Mark Warren have taken the responsibility for Cronus
> > > test, both die sort and test of packaged parts,

> > > Plan : Same as packaging : by the end of May, we want to decide
> > > on our plan so that only implementation issues are left.

> > > Action Items : Johnny and Mark own this, so they can decide.
> > > STATUS - No progress was reported.

> > Verification:

> > New Action Items : * Need to csyn/celltest custom blocks. Major missing
> > piece at this point is the Verilog. We could use
> > some logic/verification help on this one.

> > * Atlas-based designs have latches, so we need to do
> > phase checking. Hopper to work on getting 'gloss'
> > running for Atlas.

> > Regards,
> > Drew
> >

From: graham (Graham Y. Mostyn)
Sent: Monday, May 22, 1995 11:56 AM
To: cronus
Subject: Re: Summary of Cronus/Atlas Meeting on May 19

----- Begin Included Message -----

>From: graham Mon May 22 09:50:10 1995
>Date: Mon, 22 May 1995 09:50:09 -0700
>From: graham (Graham Y. Mostyn)
>To: wingard
>Subject: Re: Summary of Cronus/Atlas Meeting on May 19
>Cc: dane, bill, graham, geert
>Content-Length: 9471

Drew, I see from the summary that I/O byte is behind schedule.
Dane tells me that his schematics are completed, and that the issue is layout resource.
Should we locate a layout contractor - or have the circuit designers do some layout?

Please let me know how we can assist to regain the schedule.
Regards, Graham.

> From: wingard Sun May 21 23:06:45 1995
> Date: Sun, 21 May 1995 23:06:32 -0700
> From: wingard (Drew Wingard)
> To: cronus
> Subject: Summary of Cronus/Atlas Meeting on May 19
> Content-Length: 8502

> Hi,
>
> Here is a summary of last weeks Cronus/Atlas Status meeting
>
> Attendees: tbr, lisar, trancy, paulp, bill, stick, vo, ong, hopper, wampler,
> brianl, fwo, wingard
>

> Note: Lines beginning with ">" are from Geert's message (of 5/9) reporting the
> results of the meeting on May 5th.

> > 1. Baseplate

> >
> > Current status : We have a gards-compiled cronus baseplate
> > which can be used for routing experiments,
> >

> > Plan : By May 31, we want to have a final baseplate.
> >

> > Action items : * Decide on final die size : Drew
> > STATUS - Makefile and tools nearly in place to support this
> > activity. Should make June 1 deadline, but in
> > some danger of slipping.
> >

> > * Pading assignment : Warren
> > STATUS - On hold pending floorplanning.
> >

> > * Floorplan, custom block
> > placements : Warren
> > STATUS - On hold pending die size work. Power interface
> > (hemming) cell work is ongoing. First version of
> > clock mast cut-outs released.
> >

> > * Clock Tree generation : Kurt, Bill
 > STATUS - In progress. Bill has communicated plans for a
 > single-mast clock system to Kurt. Kurt indicates
 > on-schedule completion is very likely.

> > * Build dummy cells for
 > TTL I/O and : B.P
 > IOBYTE : Dane
 > STATUS - No progress to report.

> > 2. Custom blocks

> > Current status : GTLB done
 > Register file : end of this week
 > Caches : layout at top-level
 > Tags : layout has not started yet
 > NB : layout has just started
 > TTL I/O : not designed yet
 > IOBYTE : not started yet

> > Plan : finish all custom blocks by end of May
 > except TTL I/O and IOBYTE

> > Action Items : * Register File : B.P, Mike, Orlando
 > STATUS - Polygons finished on 5/19. Largely DRC/LVS clean.
 > Top-level should be LVS clean by 5/24.

> > * Caches : Bruce, Efelias, Dtacmo
 > STATUS - Laying out 'last' two cells. Final size is known.

> > * Tags : Bruce, Efelias, Dtacmo
 > STATUS - No progress, but can use only the cache leaf
 > cells and therefore be only top-level assembly.

> > * NB : Vikki, B.P.
 > STATUS - Polygons should be finished by about 5/24.

> > * TTL I/O : B.P., ??
 > STATUS - Design goals specified. In design. Some early
 > layout work has begun?

> > * IOBYTE layout can start by end of next week.
 > We'll need schematics by then : Dane, Bill
 > STATUS - No progress was reported.

> > New Action Items : * IOBYTE is apparently slipping badly. It will not
 > make the June 1 deadline. Drew to find out what we
 > can do to make better progress.

> > * XLU: assumption was that we would try to build XLU
 > datapath out of Atlas std cells (with 'custom'
 > placement and hand route). Tbr to ask Karzes about
 > modifying generation program to output std cells.
 > Drew to consider adding mux16 to support Stage 2.

> > * CERPOKGEN: Power OK circuit needs to be designed.
 > Since the caches will be a client, consider using
 > circuit that Bruce had started designing.

> > 3. Atlas database

> > Current status : Database builds from toplevel. An initial version
 > of about everything is there.
 > STATUS - Atlas database builds from scratch as of 5/18

> > Plan : Have a complete and accurate Atlas database by the end of May

```

> > In building Cronus sub-blocks, everybody should point to
> > /u/chip/atlas and not local versions. That will catch
> > database problems early.
> >
> > Action Items : * Review XL, XS cells
> > timing, layout, cap models : Warren
> > STATUS - All cells have passed corner analysis. Some
> > had to be resized for better margin. Latches
> > and muxes have been performance tuned. Other
> > cells ongoing.
> >
> > * Finish timing numbers : Fred
> > STATUS - No progress to report. Hopper to help with
> > capacitance and input loading work.
> >
> > * Complete verilog libraries : Brianl
> > STATUS - Basic libraries installed. Still need models
> > for Cerberus and 'SC' cells.
> >
> > * Complete Synopsis libraries : Brianl
> > STATUS - Basic libraries installed. Severe performance
> > degradation with latest version of Synopsys
> > noticed, but can still use prior version until
> > fix/workaround is provided.
> >
> > New Action Items : * Enhance Verilog libraries : Brianl and Drew
> > Tbr has provided a list of 'missing' cells
> > that prevent him from compiling Cronus for
> > simulation at the top level. We need to get
> > the truly needed ones installed ASAP.
> >
> > * Enhance Gate Family : Warren and Drew
> > We need XOR and MAJ gates, and perhaps an
> > asynchronously-loadable latch for Cerberus.
> >
> > 4. Cronus/verilog/bsrc, Makefile.rules, GARDS
> >
> > Current status : Initial versions of Makefile.rules exist. We can
> > build a sub-block using the atlas database on
> > the checked in baseplate model.
> >
> > Plan : Have a "good" Makefile.rules that works well enough that
> > other people can start mapping Euterpe blocks by the end of
> > May
> > Implement " put the top-level together" strategy in the
> > equivalent of euterpe/verilog/bsrc/Makefile.tst
> >
> > Action Items : * Makefile.rules : Brianl
> > STATUS - "Better" version released. Ongoing.
> >
> > * GARDS model additions to deal
> > with power and clock obs : Tau, Kurt
> > STATUS - Completed and tested - works great!
> >
> > * Pim2pif upgrade : Hopper
> > STATUS - Completed.
> >
> > * Take euterpe snapshot : Tbr
> > STATUS - Completed. Cronus.V released.
> > Tbr has been cutting out Euterpe blocks that
> > Cronus will not have, and bringing over
> > behavioral simulation models of the custom
> > blocks. Billz has swapped in a behavioral
> > NB buffer description. Dickson has been
> > modifying the Cerberus register map to remove
> > bits associated with unused functionality.
> >

```

```

> > * Floorplanning tools : Drew
> STATUS - Makefile and tools nearly in place to support this
> activity. Should make June 1 deadline, but in
> some danger of slipping.
>
> > * verilog/bsrc/Makefile : Drew
> STATUS - See above. Rest of this activity is largely in
> the hands of Tbr and Brianl.
>
> > 5. "Implementation of mapping"
>
> No plans for this month
> STATUS - Brianl has been 'test' mapping the top-level
> Euterpe blocks to identify mapping and P&R issues.
> This should also provide valuable feedback in
> setting the die size.
>
> > 6. Packaging
>
> Current status : We have "a plan" (read all about it in Mosaic)
>
> Plan : By the end of May, make decisions on all aspects of "the plan"
>
> Action Items : * Call a set of meetings
> during May to turn our
> plan into a set of decisions : Geert
> STATUS - Met on 5/15 and decided on 70mm TAB frame.
> Trancy investigating having someone else do
> ILB/OLB work, but lead times are about as
> bad as doing it ourselves. Drew spoke with
> MMS and got clarification on module spacing
> requirements. Trancy needs firm substrate size
> to proceed much further. Drew to provide that
> this week.
>
> > 7. Test
>
> Current Status : Mudge & Mark Warren have taken the responsibility for Cronus
> test, both die sort and test of packaged parts,
>
> Plan : Same as packaging : by the end of May, we want to decide
> on our plan so that only implementation issues are left.
>
> Action Items : Johnny and Mark own this, so they can decide.
> STATUS - No progress was reported.
>
> Verification:
> New Action Items : * Need to csyn/celltest custom blocks. Major missing
> piece at this point is the Verilog. We could use
> some logic/verification help on this one.
>
> * Atlas-based designs have latches, so we need to do
> phase checking. Hopper to work on getting 'gloss'
> running for Atlas.
>
> Regards,
> Drew
>

```

----- End Included Message -----

From: graham (Graham Y. Mostyn)
Sent: Monday, May 22, 1995 11:50 AM
To: wingard
Cc: dane; bill; graham; geert
Subject: Re: Summary of Cronus/Atlas Meeting on May 19

Drew, I see from the summary that I/O byte is behind schedule. Dane tells me that his schematics are completed, and that the issue is layout resource. Should we locate a layout contractor - or have the circuit designers do some layout?

Please let me know how we can assist to regain the schedule.
Regards, Graham.

```
> From wingard Sun May 21 23:06:45 1995
> Date: Sun, 21 May 1995 23:06:32 -0700
> From: wingard (Drew Wingard)
> To: cronus
> Subject: Summary of Cronus/Atlas Meeting on May 19
> Content-Length: 8502
>
> Hi,
>
> Here is a summary of last weeks Cronus/Atlas Status meeting
>
> Attendees: tbr, lisar, trancy, paulp, bill, stick, vo, ong, hopper, wampler,
> brianl, fwo, wingard
>
> Note: Lines beginning with ">" are from Geert's message (of 5/9) reporting the
> results of the meeting on May 5th.
>
> > 1. Baseplate
> >
> > Current status : We have a gards-compiled cronus baseplate
> > which can be used for routing experiments,
> >
> > Plan : By May 31, we want to have a final baseplate.
> >
> > Action items : * Decide on final die size : Drew
> > STATUS - Makefile and tools nearly in place to support this
> > activity. Should make June 1 deadline, but in
> > some danger of slipping.
> >
> > * Padring assignment : Warren
> > STATUS - On hold pending floorplanning.
> >
> > * Floorplan, custom block
> > placements : Warren
> > STATUS - On hold pending die size work. Power interface
> > (hemming) cell work is ongoing. First version of
> > clock mast cut-outs released.
> >
> > * Clock Tree generation : Kurt, Bill
> > STATUS - In progress. Bill has communicated plans for a
> > single-mast clock system to Kurt. Kurt indicates
> > on-schedule completion is very likely.
> >
> > * Build dummy cells for
> > TTL I/O and : B.P
> > IOBYTE : Dane
> > STATUS - No progress to report.
```

> > 2. Custom blocks

> >
> > Current status : GTLB done
> > Register file : end of this week
> > Caches : layout at top-level
> > Tags : layout has not started yet
> > NB : layout has just started
> > TTL I/O : not designed yet
> > IOBYTE : not started yet
> >
> > Plan : finish all custom blocks by end of May
> > except TTL I/O and IOBYTE
> >
> > Action Items : * Register File : B.P, Mike, Orlando
> > STATUS - Polygons finished on 5/19. Largely DRC/LVS clean.
> > Top-level should be LVS clean by 5/24.
> >
> > * Caches : Bruce, Efeliass, Dtacmo
> > STATUS - Laying out 'last' two cells. Final size is known.
> >
> > * Tags : Bruce, Efeliass, Dtacmo
> > STATUS - No progress, but can use only the cache leaf
> > cells and therefore be only top-level assembly.
> >
> > * NB : Vikki, B.P.
> > STATUS - Polygons should be finished by about 5/24.
> >
> > * TTL I/O : B.P., ??
> > STATUS - Design goals specified. In design. Some early
> > layout work has begun?
> >
> > * IOBYTE layout can start by end of next week.
> > We'll need schematics by then : Dane, Bill
> > STATUS - No progress was reported.

> > New Action Items : * IOBYTE is apparently slipping badly. It will not
> > make the June 1 deadline. Drew to find out what we
> > can do to make better progress.

> > * XLU: assumption was that we would try to build XLU
> > datapath out of Atlas std cells (with 'custom'
> > placement and hand route). Tbr to ask Karzes about
> > modifying generation program to output std cells.
> > Drew to consider adding mux16 to support Stage 2.

> > * CERPOKGEN: Power OK circuit needs to be designed.
> > Since the caches will be a client, consider using
> > circuit that Bruce had started designing.

> > 3. Atlas database

> >
> > Current status : Database builds from toplevel. An initial version
> > of about everything is there.
> > STATUS - Atlas database builds from scratch as of 5/18

> >
> > Plan : Have a complete and accurate Atlas database by the end of May
> > In building Cronus sub-blocks, everybody should point to
> > /u/chip/atlas and not local versions. That will catch
> > database problems early.

> >
> > Action Items : * Review XL, XS cells
> > timing, layout, cap models : Warren
> > STATUS - All cells have passed corner analysis. Some
> > had to be resized for better margin. Latches
> > and muxes have been performance tuned. Other
> > cells ongoing.

```

>
> > * Finish timing numbers : Fred
> STATUS - No progress to report. Hopper to help with
> capacitance and input loading work.
>
> > * Complete verilog libraries : Brianl
> STATUS - Basic libraries installed. Still need models
> for Cerberus and 'SC' cells.
>
> > * Complete Synopsis libraries : Brianl
> STATUS - Basic libraries installed. Severe performance
> degradation with latest version of Synopsys
> noticed, but can still use prior version until
> fix/workaround is provided.
>
> New Action Items : * Enhance Verilog libraries : Brianl and Drew
> Tbr has provided a list of 'missing' cells
> that prevent him from compiling Cronus for
> simulation at the top level. We need to get
> the truly needed ones installed ASAP.
>
> * Enhance Gate Family : Warren and Drew
> We need XOR and MAJ gates, and perhaps an
> asynchronously-loadable latch for Cerberus.
>
> > 4. Cronus/verilog/bsrc, Makefile.rules, GARDS
>
> Current status : Initial versions of Makefile.rules exist. We can
> build a sub-block using the atlas database on
> the checked in baseplate model.
>
> Plan : Have a "good" Makefile.rules that works well enough that
> other people can start mapping Euterpe blocks by the end of
> May
> Implement " put the top-level together" strategy in the
> equivalent of euterpe/verilog/bsrc/Makefile.tst
>
> Action Items : * Makefile.rules : Brianl
> STATUS - "Better" version released. Ongoing.
>
> * GARDS model additions to deal
> with power and clock obs : Tau, Kurt
> STATUS - Completed and tested - works great!
>
> * Pim2pif upgrade : Hopper
> STATUS - Completed.
>
> > * Take euterpe snapshot : Tbr
> STATUS - Completed. Cronus.V released.
> Tbr has been cutting out Euterpe blocks that
> Cronus will not have, and bringing over
> behavioral simulation models of the custom
> blocks. Billz has swapped in a behavioral
> NB buffer description. Dickson has been
> modifying the Cerberus register map to remove
> bits associated with unused functionality.
>
> > * Floorplanning tools : Drew
> STATUS - Makefile and tools nearly in place to support this
> activity. Should make June 1 deadline, but in
> some danger of slipping.
>
> > * verilog/bsrc/Makefile : Drew
> STATUS - See above. Rest of this activity is largely in
> the hands of Tbr and Brianl.
>
> > 5. "Implementation of mapping"

```

> >
> > No plans for this month
> STATUS - Brianl has been 'test' mapping the top-level
> Euterpe blocks to identify mapping and P&R issues.
> This should also provide valuable feedback in
> setting the die size.
> > 6. Packaging
> >
> Current status : We have "a plan" (read all about it in Mosaic)
> >
> Plan : By the end of May, make decisions on all aspects of "the plan"
> >
> Action Items : * Call a set of meetings
> during May to turn our
> plan into a set of decisions : Geert
> STATUS - Met on 5/15 and decided on 70mm TAB frame.
> Trancy investigating having someone else do
> ILB/OLB work, but lead times are about as
> bad as doing it ourselves. Drew spoke with
> MMS and got clarification on module spacing
> requirements. Trancy needs firm substrate size
> to proceed much further. Drew to provide that
> this week.
> > 7. Test
> >
> Current Status : Mudge & Mark Warren have taken the responsibility for Cronus
> test, both die sort and test of packaged parts,
> >
> Plan : Same as packaging : by the end of May, we want to decide
> on our plan so that only implementation issues are left.
> >
> Action Items : Johnny and Mark own this, so they can decide.
> STATUS - No progress was reported.
> Verification:
> New Action Items : * Need to csyn/celltest custom blocks. Major missing
> piece at this point is the Verilog. We could use
> some logic/verification help on this one.
> >
> * Atlas-based designs have latches, so we need to do
> phase checking. Hopper to work on getting 'gloss'
> running for Atlas.
> > Regards,
> Drew
>

From: wingard (Drew Wingard)
Sent: Monday, May 22, 1995 1:07 AM
To: cronus
Subject: Summary of Cronus/Atlas Meeting on May 19

Hi,

Here is a summary of last weeks Cronus/Atlas Status meeting

Attendees: tbr, lisar, trancy, paulp, bill, stick, vo, ong, hopper, wampler,
brianl, fwo, wingard

Note: Lines beginning with ">" are from Geert's message (of 5/9) reporting the results of the meeting on May 5th.

> 1. Baseplate
>
> Current status : We have a gards-compiled cronus baseplate
> which can be used for routing experiments,
>
> Plan : By May 31, we want to have a final baseplate.
>
> Action items : * Decide on final die size : Drew
STATUS - Makefile and tools nearly in place to support this
activity. Should make June 1 deadline, but in
some danger of slipping.

> * Padring assignment : Warren
STATUS - On hold pending floorplanning.

> * Floorplan, custom block
> placements : Warren
STATUS - On hold pending die size work. Power interface
(hemming) cell work is ongoing. First version of
clock mast cut-outs released.

> * Clock Tree generation : Kurt, Bill
STATUS - In progress. Bill has communicated plans for a
single-mast clock system to Kurt. Kurt indicates
on-schedule completion is very likely.

> * Build dummy cells for
> TTL I/O and : B.P
> IOBYTE : Dane
STATUS - No progress to report.

> 2. Custom blocks
>
> Current status : GTLB done
> Register file : end of this week
> Caches : layout at top-level
> Tags : layout has not started yet
> NB : layout has just started
> TTL I/O : not designed yet
> IOBYTE : not started yet

> Plan : finish all custom blocks by end of May
> except TTL I/O and IOBYTE

> Action Items : * Register File : B.P, Mike, Orlando
STATUS - Polygons finished on 5/19. Largely DRC/LVS clean.

Top-level should be LVS clean by 5/24.

> * Caches : Bruce, Efeliass, Dttacmo
STATUS - Laying out 'last' two cells. Final size is known.

> * Tags : Bruce, Efeliass, Dttacmo
STATUS - No progress, but can use only the cache leaf
cells and therefore be only top-level assembly.

> * NB : Vikki, B.P.
STATUS - Polygons should be finished by about 5/24.

> * TTL I/O : B.P., ??
STATUS - Design goals specified. In design. Some early
layout work has begun?

> * IOBYTE layout can start by end of next week.
We'll need schematics by then : Dane, Bill
STATUS - No progress was reported.

New Action Items : * IOBYTE is apparently slipping badly. It will not
make the June 1 deadline. Drew to find out what we
can do to make better progress.

* XLU: assumption was that we would try to build XLU
datapath out of Atlas std cells (with 'custom'
placement and hand route). Tbr to ask Karzes about
modifying generation program to output std cells.
Drew to consider adding mux16 to support Stage 2.

* CERPOKGEN: Power OK circuit needs to be designed.
Since the caches will be a client, consider using
circuit that Bruce had started designing.

> 3. Atlas database

> Current status : Database builds from toplevel. An initial version
of about everything is there.
STATUS - Atlas database builds from scratch as of 5/18

> Plan : Have a complete and accurate Atlas database by the end of May
In building Cronus sub-blocks, everybody should point to
/u/chip/atlas and not local versions. That will catch
database problems early.

> Action Items : * Review XL, XS cells
timing, layout, cap models : Warren
STATUS - All cells have passed corner analysis. Some
had to be resized for better margin. Latches
and muxes have been performance tuned. Other
cells ongoing.

> * Finish timing numbers : Fred
STATUS - No progress to report. Hopper to help with
capacitance and input loading work.

> * Complete verilog libraries : Brianl
STATUS - Basic libraries installed. Still need models
for Cerberus and 'SC' cells.

> * Complete Synopsis libraries : Brianl
STATUS - Basic libraries installed. Severe performance
degradation with latest version of Synopsis
noticed, but can still use prior version until
fix/workaround is provided.

New Action Items : * Enhance Verilog libraries : Brianl and Drew
Tbr has provided a list of 'missing' cells
that prevent him from compiling Cronus for
simulation at the top level. We need to get
the truly needed ones installed ASAP.

* Enhance Gate Family : Warren and Drew
We need XOR and MAJ gates, and perhaps an
asynchronously-loadable latch for Cerberus.

> 4. Cronus/verilog/bsrc, Makefile.rules, GARDS

>
> Current status : Initial versions of Makefile.rules exist. We can
> build a sub-block using the atlas database on
> the checked in baseplate model.

> Plan : Have a "good" Makefile.rules that works well enough that
> other people can start mapping Euterpe blocks by the end of
> May

> Implement " put the top-level together" strategy in the
> equivalent of euterpe/verilog/bsrc/Makefile.tst

> Action Items : * Makefile.rules : Brianl
STATUS - "Better" version released. Ongoing.

> * GARDS model additions to deal
> with power and clock obs : Tau, Kurt
STATUS - Completed and tested - works great!

> * Pim2pif upgrade : Hopper
STATUS - Completed.

> * Take euterpe snapshot : Tbr
STATUS - Completed. Cronus.V released.
Tbr has been cutting out Euterpe blocks that
Cronus will not have, and bringing over
behavioral simulation models of the custom
blocks. Billz has swapped in a behavioral
NB buffer description. Dickson has been
modifying the Cerberus register map to remove
bits associated with unused functionality.

> * Floorplanning tools : Drew
STATUS - Makefile and tools nearly in place to support this
activity. Should make June 1 deadline, but in
some danger of slipping.

> * verilog/bsrc/Makefile : Drew
STATUS - See above. Rest of this activity is largely in
the hands of Tbr and Brianl.

> 5. "Implementation of mapping"

>
> No plans for this month
STATUS - Brianl has been 'test' mapping the top-level
Euterpe blocks to identify mapping and P&R issues.
This should also provide valuable feedback in
setting the die size.

> 6. Packaging

>
> Current status : We have "a plan" (read all about it in Mosaic)

> Plan : By the end of May, make decisions on all aspects of "the plan"

> Action Items : * Call a set of meetings
> during May to turn our

> plan into a set of decisions : Geert
STATUS - Met on 5/15 and decided on 70mm TAB frame.
Trancy investigating having someone else do
ILB/OLB work, but lead times are about as
bad as doing it ourselves. Drew spoke with
MMS and got clarification on module spacing
requirements. Trancy needs firm substrate size
to proceed much further. Drew to provide that
this week.

> 7. Test

>
> Current Status : Mudge & Mark Warren have taken the responsibility for Cronus
> test, both die sort and test of packaged parts,
>
> Plan : Same as packaging : by the end of May, we want to decide
> on our plan so that only implementation issues are left.
>
> Action Items : Johnny and Mark own this, so they can decide.
STATUS - No progress was reported.

Verification:

New Action Items : * Need to csyn/celltest custom blocks. Major missing
piece at this point is the Verilog. We could use
some logic/verification help on this one.

* Atlas-based designs have latches, so we need to do
phase checking. Hopper to work on getting 'gloss'
running for Atlas.

Regards,
Drew

From: hopper (Mark Hofmann)
Sent: Saturday, May 20, 1995 4:36 PM
To: Geert Rosseel
Cc: cadettes; paulp (Paul Poenisch); tbr (Tim B. Robinson)
Subject: Re: Pads

Geert Rosseel writes:

Well, maybe it's because of my lack of sleep or something, but I have some trouble with this approach.

I think we should make it very clear that if a chip passes the DRC rule check, then it should be a manufacturable, yieldable and production worthy chip.

Sorry, if I wasn't clear. Yes that's what I meant and that's my understanding.

If the "guidelines" will bring the yield from 90% to 95%, well, I have no problem with that.

On the other hand, if the "guidelines" are necessary to bring the yield up from .01% (like a couple of prototypes) to a production worthy die, then I have a real problem with this.

I don't want to be in a position where we have to re-layout every design at least twice : one that follows the DRC rules and gets us a prototype and the sub-sequent re-layouts to fix up all the "guideline" violations that we only see after processing the die.

The pads have been an area where there has been much re-work. If the pads still violate the current DRC rules, then they'll have to be changed again. If the design rules have to change to guarantee a production die, then they'll have to change. I think we have a stable set of DRC rules. Now I'm not so sure about the pad layouts.

The problem has been to decide which rules are "rules" and which rules are "guidelines". There will always be an area open to interpretation, but we need to establish a consistent set which 1. designer's must follow 2. CAD must check 3. Fab will guarantee to a minimum yield. Since we don't have yield data, the third is unsettled.

I want to be in the same position as with CSM or TSMC : you follow the rules and they build a die that , if it is electrically and logically correct, can go immediately in production without a redesign.

That is CAD's intent, at least.

Finally, "guidelines" only should affect back-end CAD. It doesn't make sense to tell a mask-designer that his layout meets DRC rules, but it's no good. If anybody wants to have "guidelines" that affect manual layout, then they should volunteer the resources to do the visual inspection and sign-off on any layout as fast as a DRC can be run.

The point is that there are good and bad designs. A design using all minimum design rule widths and spacings should be correct and yield, but I would expect a design which has greater than minimum widths and spacings, that has redundant vias and so forth would yield better.

The style guide will make the point that if you do not need to use minimum width it's better not to use it. If you can put in a second via, it's a good thing. This is art. The DR checker checks for spelling errors.

If it's wrong, it's wrong. The style checker checks for wording- the design may be clumsy

or awkward but still work, yet it could be better.
Therefore, style guidelines do affect custom layout.

Most places do not codify such documents, let alone check them with a CAD tool. This is something that should raise our yield from production- worthy to great. If we find something in the style guide that is a production-stopper then it needs to be made checkable and migrate to the rule book. Likewise, if we find a rule that has less effect on yield it can be moved to the style guide, or eliminated all together.

In my opinion, at this stage of Euterpe and at this stage of our Fab, we should concentrate on the rule book and not pay too much attention to the style guide.

Geert

-hopper

From: hopper (Mark Hofmann)
Sent: Saturday, May 20, 1995 4:36 PM
To: Geert Rosseel
Cc: cadettes; paulp (Paul Poenisch); tbr (Tim B. Robinson)
Subject: Re: Pads

Geert Rosseel writes:

Well, maybe it's because of my lack of sleep or something, but I have some trouble with this approach.

I think we should make it very clear that if a chip passes the DRC rule check, then it should be a manufacturable, yieldable and production worthy chip.

Sorry, if I wasn't clear. Yes that's what I meant and that's my understanding.

If the "guidelines" will bring the yield from 90% to 95%, well, I have no problem with that.

On the other hand, if the "guidelines" are necessary to bring the yield up from .01% (like a couple of prototypes) to a production worthy die, then I have a real problem with this.

I don't want to be in a position where we have to re-layout every design at least twice : one that follows the DRC rules and gets us a prototype and the sub-subsequent re-layouts to fix up all the "guideline" violations that we only see after processing the die.

The pads have been an area where there has been much re-work. If the pads still violate the current DRC rules, then they'll have to be changed again. If the design rules have to change to guarantee a production die, then they'll have to change. I think we have a stable set of DRC rules. Now I'm not so sure about the pad layouts.

The problem has been to decide which rules are "rules" and which rules are "guidelines". There will always be an area open to interpretation, but we need to establish a consistent set which 1. designer's must follow 2. CAD must check 3. Fab will guarantee to a minimum yield. Since we don't have yield data, the third is unsettled.

I want to be in the same position as with CSM or TSMC : you follow the rules and they build a die that , if it is electrically and logically correct, can go immediately in production without a redesign.

That is CAD's intent, at least.

Finally, "guidelines" only should affect back-end CAD. It doesn't make sense to tell a mask-designer that his layout meets DRC rules, but it's no good. If anybody wants to have "guidelines" that affect manual layout, then they should volunteer the resources to do the visual inspection and sign-off on any layout as fast as a DRC can be run.

The point is that there are good and bad designs. A design using all minimum design rule widths and spacings should be correct and yield, but I would expect a design which has greater than minimum widths and spacings, that has redundant vias and so forth would yield better.

The style guide will make the point that if you do not need to use minimum width it's better not to use it. If you can put in a second via, it's a good thing. This is art. The DR checker checks for spelling errors.

If it's wrong, it's wrong. The style checker checks for wording- the design may be clumsy

or awkward but still work, yet it could be better.
Therefore, style guidelines do affect custom layout.

Most places do not codify such documents, let alone check them with a CAD tool. This is something that should raise our yield from production- worthy to great. If we find something in the style guide that is a production-stopper then it needs to be made checkable and migrate to the rule book. Likewise, if we find a rule that has less effect on yield it can be moved to the style guide, or eliminated all together.

In my opinion, at this stage of Euterpe and at this stage of our Fab, we should concentrate on the rule book and not pay too much attention to the style guide.

Geert

-hopper

From: hopper (Mark Hofmann)
Sent: Saturday, May 20, 1995 10:47 AM
To: geert (Geert Rosseel)
Cc: paulp (Paul Poenisch); cadettes; tbr (Tim B. Robinson)
Subject: Re: Pads

vant writes:

Geert,

The types of rules that Paul is referring too are 'probabalistic failure mechanisms'. Which means that in certain cases, it might fail but then, it might not.

I agree with you that unless it is checkable, then we can't ensure the chip to work.

Paul is still working on a list of things which are recomendations and putting them in a 'style guide'. I only know of one rule right now that is uncheckable, and that is 'fat metal spacing'.

This rule is complex and causing the metal synthesis to increase drc run times by about 3x, and then backend tapeout to increase by many days.

I will be out most of the weekend, so I don't know how much I can help, but I'll try.

Thanks,
Dave

Geert,

To ampilfy this: It has been agreed that there are to be 2 documents:
1. DRC Bible 2. Style Guide. All circuits must pass all rules in the the DRC Bible. For a rule to be placed in the DRC Bible it must be checkable.
There are many steps which Al and others can not give hard and fast rules for, but which they feel if done a certain way, will increase yield.
These things are to be placed in a separate Style Guide. Paul is working on this. The CAD group is going to first focus on getting the DRC Bible fully and correctly implemented. We will work on the style guide second. (A Style Guide example: Uniform spacing of certain vias. How do you define "uniform"?
When is there a violation?. We will need to implement certain area density functions to check this.)

It is my understanding- and Paul please correct me if I err- that all circuits which pass the full DRC Bible ruleset should yield functioning die, to the best of the processing group's knowledge, without the need for any further "style" or "visual" checks.

-hopper

From: tbr (Tim B. Robinson)
Sent: Saturday, May 20, 1995 12:47 PM
To: geert (Geert Rosseel)
Cc: cadettes; mikew; mudge@microunity.com; paulp
Subject: Re: Pads

Geert Rosseel wrote (on Sat May 20):

I don't know who was involved in this decision but I am very strongly opposed to having design rules which, if not met, cause the chip to fail, but at the same time are not checked.

As far as I am concerned, rules that are not written down in the design rule document and are not checked by the DRC flow. don't exist.

How do we guarantee that no similar structures exist on the die, how do we check that no similar structures exist on Pollux, Castor or any other die ? Are we committing ourselves now to visual inspection of all polygons in all future tape-outs ???

If the current rules are uncheckable, then, they have to be changed until they can be checked.

I agree on this. A rule is not a rule if it's not checked. If it's important enough to have a material effect on the result, it has to be checked 100%.

Tim

From: tbr
Sent: Saturday, May 20, 1995 12:47 PM
To: geert (Geert Rosseel)
Cc: cadettes; mikew; mudge@microunity.com; paulp
Subject: Re: Pads

Geert Rosseel wrote (on Sat May 20):

I don't know who was involved in this decision but I am very strongly opposed to having design rules which, if not met, cause the chip to fail, but at the same time are not checked.

As far as I am concerned, rules that are not written down in the design rule document and are not checked by the DRC flow, don't exist.

How do we guarantee that no similar structures exist on the die, how do we check that no similar structures exist on Pollux, Castor or any other die ? Are we committing ourselves now to visual inspection of all polygons in all future tape-outs ????

If the current rules are uncheckable, then, they have to be changed until they can be checked.

I agree on this. A rule is not a rule if it's not checked.

If it's important enough to have a material effect on the result, it has to be checked 100%.

Tim

From: vanthof (vant)
Sent: Saturday, May 20, 1995 12:32 PM
To: Geert Rosseel
Cc: vanthof (Dave Van't Hof); hopper (Mark Hofmann)
Subject: Re: Pads

Geert Rosseel writes:

>
> I don't know who was involved in this decision but I am very strongly
>opposed to having design rules which, if not met, cause the chip to
>fail, but at the same time are not checked.
>
> As far as I am concerned, rules that are not written down in the
>design rule document and are not checked by the DRC flow.
>don't exist.
>
> How do we guarantee that no similar structures exist on the die, how
>do we check that no similar structures exist on Pollux, Castor or any
>other die ? Are we committing ourselves now to visual inspection of all
>polygons in all future tape-outs ???
>
> If the current rules are uncheckable, then, they have to be changed
>until they can be checked.
>
> Geert

Geert,

The types of rules that Paul is referring too are 'probabalistic failure mechanisms'. Which means that in certain cases, it might fail but then, it might not.

I agree with you that unless it is checkable, then we can't ensure the chip to work.

Paul is still working on a list of things which are recommendations and putting them in a 'style guide'. I only know of one rule right now that is uncheckable, and that is 'fat metal spacing'. This rule is complex and causing the metal synthesis to increase drc run times by about 3x, and then backend tapeout to increase by many days.

I will be out most of the weekend, so I don't know how much I can help, but I'll try.

Thanks,
Dave

P.S. How is Mathias and Leiva? (I'm a horrible speller, sorry) Got any pictures to show off yet???

--
Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100
"Don't blame me! I didn't vote for him"

From: geert (Geert Rosseel)
Sent: Saturday, May 20, 1995 12:20 PM
To: mudge@microunity.com; paulp
Cc: cadettes; mikew; tbr
Subject: Re: Pads

Johnny said :

>>
>> Geert,
>> Mike was in my office today saying that the supply switch on the pads
>> is done and they pass lvs. Well done Mike! The next step I believe
>> is to incorporate the new design rules or guidelines. We need input
>> from Paul to make that happen and preferably some more of Mike's
>> time.
>> Bug me about this when you get back.
>>
>> johnnymudge
>>
>>
>>

And then Paulp says :

> Johnny,
>
> I was hoping to have an unofficial version of the design rules to the
> mask designers so these changes could be made last Thursday.
> Unfortunately things were still in flux. It looks like I may be able to get it out
tomorrow.
>
>
>
> However I will need to work some with whomever is going to make these
> changes (Mike I assume) because we have decided that changes in the
> design rules to eliminate all the problems we have with these cells
> would be uncheckable. This means that the changes will be based on
> verbal or written descriptions of what needs to be done rather than on meeting DRC's.

I don't know who was involved in this decision but I am very strongly opposed to having design rules which, if not met, cause the chip to fail, but at the same time are not checked.

As far as I am concerned, rules that are not written down in the design rule document and are not checked by the DRC flow don't exist.

How do we guarantee that no similar structures exist on the die, how do we check that no similar structures exist on Pollux, Castor or any other die ? Are we committing ourselves now to visual inspection of all polygons in all future tape-outs ???

If the current rules are uncheckable, then, they have to be changed until they can be checked.

Geert

From: Curtis Abbott [abbott@microunity.com]
Sent: Thursday, May 18, 1995 4:37 PM
To: Mark Semmelmeier
Cc: billz@microunity.com; craig@microunity.com; dickson@microunity.com; dit00@microunity.com; gmo@microunity.com; guarino@microunity.com; hayes@microunity.com; jeffm@microunity.com; khp@microunity.com; lisar@microunity.com; puri@microunity.com; tbr@microunity.com; woody@microunity.com; yam@microunity.com
Subject: GGFMul restartability bug

I should think

1. we could easily teach the compiler not to generate a ggfmul8 with src reg == dest reg (yes, Ray?) 2. this would cause the bug to not matter (yes, Mark?)

If so, it would be our first software workaround of a hardware bug, implemented before tapeout! I can't imagine that this workaround would impact performance in any significant way.

- Curtis

From: mws (Mark Semmelmeyer)
Sent: Thursday, May 18, 1995 2:25 PM
To: abbot; craig; gmo; guarino; hayes; khp; puri; tbr; yam
Cc: billz; dickson; dit00; jeffm; lisar; woody
Subject: GGFMul restartability bug

Until yesterday (bsrc BOM 306.0), GGFMul had the following bug:
If a source register pair had a pending non-blocking load to only its high octlet and was not the same register pair as the destination register pair, then the cylinder issuing the GGFMul would hang.
The hung cylinder would indirectly hang NB and ignore interrupts.

I had a difficult time trying to fix this. Finally I thought that I should just allow the GGFMul to become nonatomic halfway through so that the nonblocking load return could get through. This had the side effect of allowing the instruction to be interrupted halfway through, which we had already accepted in certain (rarer) hiccup cases. This fix was very cheap and caused almost no disturbance to the near-tapeout design (bsrc BOM 307.0).

Today I realized that the fix damages the case of a source register shared with the destination register. Since now GGFMul is generally interruptible and modifies its destination in 2 separate octlet writes, a restart would see a modified source. I currently believe this is only possible when the GGFMul issues on etal, and JeffM is going to try to tickle it. Larry Yamano showed me Brendan's ?rs_syndrome.c? code and it was doing a GGFMul, then an XOR, and then feeding that back into the next loop iteration's GGFMul. This looks like does/could avoid GGFMul with src=dst. Abbott may know of any other GGFMul code, but it sounded like there were very few cases expected. Of course, if the GGFMul cylinder is not enabled for interrupts, this is all a moot issue.

I might still be able to fix this by adding 15 xors, maybe 5 flops of buffering, and 3-4 orflops to Issue to patch this up. But maybe we want to avoid the disturbance at this late date. Can software survive with the less-than-perfect implementation of BOM 307?

From: wampler (Kurt Wampler)
Sent: Thursday, May 18, 1995 12:14 PM
To: ken
Cc: hopper; sysadmin; tom; vanthof
Subject: Low-level format of trex /s1

Hi, Ken -

Now is a good time to reformat the /s1 3-spindle volume set on trex in preparation for euterpe tapeout, ***provided that this can be done without having to reboot trex.*** Dave has an important metal synthesis test running on trex right now that we don't want to interrupt, and it may continue for several days.

What I think needs to be done is:

- 1) Back up the contents of trex /s1 on tape.
- 2) Reformat the spindle containing the bad file (or reformat all 3 spindles if the particular spindle cannot be identified).
- 3) Re-initialize the /s1 volume set and restore the contents from the backup tape.

Thanks,

Kurt

From: hopper (Mark Hofmann)
Sent: Thursday, May 18, 1995 2:20 AM
To: Tim B. Robinson
Subject: Re: 10am AI meeting

Tim B. Robinson writes:

Was there a meeting with al this morning? If so what was the outcome?
Did mouss attend?

Tim

There was a meeting. Mouss did not attend.
I would say AI appeared much more muted this time. He kept saying we should go after the "low hanging fruit" and not spend time trying to fix the last 1% of yield.

It was decided that many things we were working on should go in a separate "style guide" document. This document will basically say things like, "If you don't have to use min spaced wires or min transistors, then don't do it."
However, you can violate this "style" and still get a working, but low-yielding chip. Many of the style guide rules will be difficult to check ("uniform density of vias", for example- how uniform is that?)

The Design Rule document, on the other hand, will contain rules which can be checked and which can't be violated. No new rules were added to this document in this meeting and some were relaxed and moved to the style guide.
The result was that Paul had some design rules which he did not pass out at the meeting. Instead he revised them again and passed out copies to the various CAD folks after the meeting. Dave and Tom went over these- the main issue was treatment of fat metals- and I believe have reached agreement with Paul on how to proceed. [Basically, fat metals are allowed up to 4.5u and will be post-processed. Dave and Tom believe they have a post-processing solution which will be DRC clean. It will add run time. How much we don't know yet. Dave has a test running on tomato. (Mothra crashed with a bad SIMM, vancleef has been notified). Alternatively we could restrict fat metals to 2.5u and not do post-processing. I think this would make the analog folks unhappy and present power routing obstacles.]

Paul will work on a separate style guide, but the CAD group will focus on the Design rule document.

I said that the logicians may be presenting us with a fully routed, timed Euterpe within a week, so that we should get ready to begin fracture and get masks out. The first 14 or so masks so fracture very quickly (within a week, I believe).

-hopper

From: tbr
Sent: Wednesday, May 17, 1995 11:57 PM
To: wampler (Kurt Wampler)
Cc: geert; hopper; wampler
Subject: GARDS 7.117?

Kurt Wampler wrote (on Tue May 9):

Last Friday we received a fixed GPLACE which closes out the two TARs we had registered against it. I have tested it with the original cases we submitted on the bug tape and it worked correctly. SVR will be pressing for us to accept the 7.117 release soon so they can make the source escrow of it for us.

When would be a good time to switch over to the 7.117 release for general use? There may still be other bugs lurking in this release that we haven't yet uncovered, but it has received a moderate amount of exercise by Brian, Drew, and me.

The PGRROUTE fix is still a week or two away. It would be good to defer the capture of the escrow tape until it includes the fixed PGRROUTE.

Comments?

Have we made this switch? I hope to get another top level run going tomorrow, and it might be a good test case.

Tim

From: doi (Derek Iverson)
Sent: Wednesday, May 17, 1995 1:25 PM
To: compiler; mediacom; warren; dit00; guarino; gmo; lisa; iimura; gregg; deborah; claseman
Cc: billz; mws; woody; dickson; tbr
Subject: Expanded SW BU Meeting: SW Debugging in Gate Level Environment

Today's 'Software Bringup Meeting' is going to be held in the 'Multimedia Demo Room' and will cover

Topics in Software Debugging in a Gate Level Environment

This session may well exceed the normally scheduled 1 hour.

Attendees are expected to be familiar with the contents of Euterpe Microarchitecture book.

The agenda will be something like....

1. Discussion of the simulation environment.
2. Discussion of the logfiles from the hardware simulators.
3. Discussion of the likedriverlog (gate level trace).
4. Common failure signatures.

BYOCMD (Bring your own chair and microarchitecture document).

doi

From: trancy [trancy@charybdis]
Sent: Tuesday, May 16, 1995 4:49 PM
To: Bill Herndon; Drew Wingard; Geert Rosseel; Herman Chu; Paul Poenisch; Tim B. Robinson; steve@charybdis; trancy@charybdis
Subject: FW: Cronus TAB Frame Meeting Minutes

From: trancy on Tue, May 16, 1995 11:06 AM
Subject: RE: Cronus TAB Frame Meeting Minutes
To: john mudge

I assume we will mount the TABed device straight onto the product board as Calliope and Euterpe.

I talked with Kevin of Yamaichi and he told me they do have the standard test socket for 70mm TAB frame, but I would like to review the actual drawing to confirm that.

Regards.
Trancy.

From: john mudge on Tue, May 16, 1995 11:00 AM
Subject: RE: Cronus TAB Frame Meeting Minutes
To: trancy
Cc: al; steve; bill@gaea; wingard@gaea; fwo@gaea; geert@gaea; hchu@gaea; warren@gaea; paulp@gaea; tbr@gaea; mudge

Trancy,

> I was concerned about the test socket for 70mm and was told that we
> will
100%
> test on the wafer/die level, not the TAB frame.

Yes, that was one approach that was proposed and can work. It does mean that it will be a bit more difficult to monitor M.M.S.'s flip chip bonding yield. Do we plan on mounting the TABed device straight onto the product board or will there be an intermediate daughter board?

I just requested Yamaichi to
> send/fax me some socket info. for the 70mm TAB frame, will show you as
> soon as I receive it.

Do you have knowledge that there is such a device?

johnnymudge

From: wingard (Drew Wingard)
Sent: Monday, May 15, 1995 8:22 PM
To: microlunatics
Subject: This week's SITN Seminars

Here are the abstracts for the EE 310, EE 380, and CS 547 Seminars this week. The CS 547 Seminar is new to this list: it focuses on Human-Computer Interaction. Thanks to Paul Berry for giving me the pointer to the on-line abstracts for this class: <http://www-pcd.stanford.edu/seminar.html>

Seminar	Channel	Day	Time	Live?
EE 310	E2	Tues	5:45-6:35pm	Taped
EE 380	E3	Thurs	8:00-9:15am	Taped
CS 547	E2	Fri	12:30-2:00pm	Live

Regards,
Drew

EE 310 Seminar (tomorrow)
on
Chemical-Mechanical Polishing Process for Dielectric Planarization

Milind Weling
VLSI Technology, Inc.

Tuesday, May 16, 1995, 4:15 pm, McCullough 134

Shrinking metal linewidths and and inter-metal spaces have posed several challenges for inter-metal dielectric planarity. Reduced inter-metal spaces simultaneously require improved gap fill capabilities of the inter-metal dielectric and better local planarization. Reduced minimum features require higher resolution photolithography for their patterning at metal and via levels. The concomitant reduction in the available depth of focus thereby demands improvement in the level of global planarization to maintain critical dimension (CD) control of all features over the die. In the case of typical 0.35um technologies, global planarization with topography variations less than 0.5um over the die is required.

Chemical-mechanical polishing (CMP) with its significantly improved global planarity and reduced thermal budget is clearly the leading process solution in this area. CMP allows preferential removal of material in elevated areas thereby planarizing topography over several millimeters. This talk will review the limitations of conventional dielectric planarization techniques and show how CMP improves long range planarity enabling multiple (4-5) levels of metallization. CMP process development and the trade-offs involved in the choice of consumables such as polishing slurry, pads, carrier film will be reviewed. Data from electrical, SEM and AFM characterization of the CMP process will also be reviewed. Other rapidly evolving semiconductor processing applications for the CMP process such as damascene, tungsten plug polishing will also be briefly reviewed.

Biography

Milind Weling is a Senior Process Development Engineer in the Technology Development group at VLSI Technology, Inc. He is currently responsible for developing chemical-mechanical polishing and plasma etch processes for submicron CMOS. He received the B.Tech degree in electrical engineering from the Indian Institute of Technology, Bombay in 1988. He later received the M.S. degree in electrical engineering from the University of Hawaii at Manoa in 1990. He has co-authored over 35 papers, mostly in plasma etch and CMP, and holds 3 patents.

EE380 Computer Systems Colloquium

Spring Quarter 1994-1995

Lecture #7

Date: Wednesday, May 17, 1994

Time: 4:15-5:30 pm

Location: Terman Auditorium

SITN: Thursday, Channel E3, 8:00-9:15

Speaker: Robert Miller
Producer-Director
Pure Grain Digital Productions
Palo Alto, California

Title: A CASE STUDY IN DIGITAL FILM PRODUCTION

ABSTRACT

Traditional live-action film production methods are prohibitively expensive and creatively constraining for many independent filmmakers. Computer technology has the potential to change this. High-resolution digital cameras can be used in conjunction with software running on affordable desk-top PCs to deliver the visual quality of film at a small fraction of the cost.

This talk will discuss the economic and creative advantages achieved on "Mail Bonding," the first live-action movie to exploit digital technology at every stage. To explore creative possibilities and reduce the likelihood of on-set surprises, pre-production elements such as sets, models, and storyboards were developed on desktop computers. The production was then shot with a special wide-screen digital camera on loan from Sony and recorded direct to disk. Mac-based digital playback and editing software were used extensively on the set for real-time scene evaluation and to assure continuity, and in post-production for editing, image compositing, and other visual effects.

Eventually, a digital master was created and scanned to 35mm film. The 12-minute comedy short, which will be presented in its entirety during this talk, has enjoyed a good reception at film festivals in New York, Los Angeles, Chicago, Dallas, Montreux, and Park City (Sundance).

BIOGRAPHY

Robert Miller's background teeters between writing, filmmaking, computer science, and outright wanderlust. His experiences include oil exploration in Saudi Arabia, an education in Computer Science, photo-journalism in Australia, and video producing and multimedia coordinating for the Stanford Computer Forum and Apple Computer. Robert is currently touring film festivals with his short film, "Mail Bonding," and is developing a feature length film that he will direct and produce.

People, Computers, and Design Seminar

KidSim: End User Programming of Simulations

Allen Cypher and David C. Smith, Apple

Stanford University May 19, 1995

KidSim is an environment that allows children to create their own simulations and video games. They create their own characters, and they create rules that specify how the characters are to behave and interact. KidSim is programmed by demonstration, so that users do not need to learn a conventional programming language or scripting language. Informal user studies have shown that children are able to create simulations in KidSim with a minimum of instruction, and that KidSim stimulates their imagination.

We will demonstrate the system, discuss its design, and show some worlds that children have built.

Allen Cypher is a Research Scientist at Apple Computer, Inc. His main interest is in end-user programming -- giving all computer users capabilities that have traditionally belonged to programmers. Allen is a co-inventor of KidSim. He also created the Eager system, which observes a users' actions and creates programs to automate repetitive activities. He edited the book "Watch What I Do: Programming by Demonstration", published in 1993 by MIT Press.

While on the Xerox "Star" project, David Canfield Smith invented icons and the desktop metaphor for computer interfaces, today used by 100 million people. For the past several years at Apple, Dave has worked on educational software, particularly on finding a way for children to program computers. KidSim(tm) is the culmination of that work. Dave's research interests include human-computer interaction, educational software, programming language design, programming environments, end-user programming, and getting rid of the "priests" of computing. The unifying theme behind his work for the past twenty years has been the attempt to make computers more accessible to ordinary people.

From: graham (Graham Y. Mostyn)
Sent: Monday, May 15, 1995 2:30 PM
To: andrew; rich; dane; yves; arya; rmm; geert
Cc: pollux
Subject: Re: Revision 2 of Pollux

Unfortunately my attendance at this meeting has been pre-empted by a teleconference with Cox from 1.30 - 2.30. However, I suggest that you all go ahead as planned.

Graham.

> From: graham Fri May 12 15:08:35 1995
> Date: Fri, 12 May 1995 15:08:15 -0700
> From: graham (Graham Y. Mostyn)
> To: andrew, rich, dane, yves, arya, rmm, geert
> Subject: Revision 2 of Pollux
> Cc: graham, pollux
> Content-Length: 471
>
> Test engineering has asked us how revision 2 of Pollux will differ
> from revision 1, and whether some modules might be exchanged for
> others. (For example, rev.1 contains two versions of the audio DAC,
> but the audio ADC had not yet been designed at the time of tapeout.)
>
> We will meet at 2pm on Monday, in the hardware engineering room, to
> discuss:
> - the revision differences
> - what module exchanges might be warranted
> - an update on the Callioped test results.
>
> Graham.
>

From: graham (Graham Y. Mostyn)
Sent: Friday, May 12, 1995 5:08 PM
To: andrew; rich; dane; yves; arya; rmm; geert
Cc: graham; pollux
Subject: Revision 2 of Pollux

Test engineering has asked us how revision 2 of Pollux will differ from revision 1, and whether some modules might be exchanged for others. (For example, rev.1 contains two versions of the audio DAC, but the audio ADC had not yet been designed at the time of tapeout.)

We will meet at 2pm on Monday, in the hardware engineering room, to discuss:

- the revision differences
- what module exchanges might be warranted
- an update on the Callioped test results.

Graham.

From: chip (Potatoe Chip)
Sent: Friday, May 12, 1995 4:01 PM
To: geert
Subject: output of pollux/checkoutrc

Fri May 12 12:55:27 PDT 1995 (geert Fri, 12 May 1995 12:55:16 -0700) pollux
[Release BOM (V7.0) in pollux (Fri May 12 12:55:27 PDT 1995)]

Dir	pollux	BOM 7.0
1.7	.checkoutrc	
1.6	Makefile	
6.4	welcome.ht	(No)
Dir	pollux/baseplate	BOM 8.0
1.1	.checkoutrc	
1.28	Makefile	(1.25)
1.6	baseplate.sgen.m4	(1.5)
1.2	chipparms.m4	
1.3	clean_request	(1.1)
1.2	floorplan.sgen.m4	
1.6	floorplanmod1.sgen	
1.4	floorplanmod10.sgen	
1.4	floorplanmod11.sgen	
1.5	floorplanmod12.sgen	
1.4	floorplanmod2.sgen	
1.4	floorplanmod3.sgen	
1.4	floorplanmod4.sgen	
1.5	floorplanmod5.sgen	
1.5	floorplanmod6.sgen	
1.2	floorplanmod7.sgen	
1.3	floorplanmod8.sgen	
1.2	floorplanmod9.sgen	
3.2	knobmap.ly.orig	
7.2	mod1.list	(No)
7.1	mod2.list	(No)
7.2	pad.awk	(No)
1.3	padlabels.sgen.m4	
1.3	padmodule.sgen.m4	
1.16	testmod1.sgen.m4	(1.14)
1.17	testmod10.sgen.m4	
1.10	testmod11.sgen.m4	
1.16	testmod12.sgen.m4	
1.11	testmod2.sgen.m4	(1.10)
1.15	testmod3.sgen.m4	(1.14)
1.12	testmod4.sgen.m4	(1.11)
1.13	testmod5.sgen.m4	(1.12)
1.16	testmod6.sgen.m4	
1.8	testmod7.sgen.m4	
1.16	testmod8.sgen.m4	
1.11	testmod9.sgen.m4	
1.3	testtempl.m4	
Dir	pollux/compass	BOM 4.0
1.7	vlsi.booc-all	(1.6)
1.7	vlsi.booc-dcell	(1.6)
1.6	vlsi.booc-tapeout	(1.5)
Dir	pollux/dcell	BOM 4.0
1.1	.checkoutrc	
1.2	Makefile	
Dir	pollux/doc	BOM 4.0

3.1	overview.ht	(No)
3.1	pollux.ht	(No)
Dir	pollux/gards	BOM 4.0
1.1	.checkoutrc	
1.12	Makefile	
1.6	sclean-request	
1.3	sofa_model	
Dir	pollux/ged	BOM 5.0
1.1	.checkoutrc	
1.8	Makefile	(1.7)
1.1	README	
Dir	pollux/misc	BOM 3.0
1.1	gards.ctrl	
1.1	gards.spn	
1.2	gards.vrf	
1.1	global.net	
Dir	pollux/verify	BOM 6.0
4.1	.checkoutrc	
4.1	Makefile	
Dir	pollux/verify/standalone	BOM 7.0
4.1	.checkoutrc	
4.1	Makefile	
1.2	template	
Dir	pollux/verify/standalone/mod10	BOM 7.0
6.1	.checkoutrc	
1.5	Makefile	
3.1	doc	
6.1	ios	
1.2	mod10.in	
1.2	mod10.out	
1.3	mod10_driver.V	
1.1	ram.ref	
Dir	pollux/verify/standalone/mod11	BOM 7.0
5.1	.checkoutrc	
1.6	Makefile	
1.7	check_mod.pl	
1.2	del_line	
1.1	ios	
1.1	mod11_8.sen	
1.10	tst_mod11.pl	
Dir	pollux/verify/standalone/mod12	BOM 7.0
4.1	.checkoutrc	
1.8	Makefile	
1.3	check_mod.pl	
1.1	del_line	
1.3	ios	
1.6	tst_mod12.pl	
Dir	pollux/verify/standalone/mod3	BOM 4.0
3.1	Makefile	
1.1	ios	
Dir	pollux/verify/standalone/mod8	BOM 9.0
5.1	.checkoutrc	
1.6	Makefile	
6.1	clean-request	
1.4	ios	
1.1	spy.h	
1.9	tester.V	

1.3	tester.h	
1.1	tlb.h	
1.4	tlblogical.V	
1.2	tlbmodel.V	
Dir	pollux/verify/standalone/mod9	BOM 3.0
2.1	Makefile	
1.1	mod9.srl	
1.1	mod9_196608.sen	
1.1	mod9_32768.sen	
1.1	mod9_4096.sen	
1.1	mod9_512.sen	
1.1	mod9_64.sen	
Dir	pollux/verilog	BOM 6.0
1.1	.checkoutrc	
1.3	Makefile	
Dir	pollux/verilog/lvs	BOM 5.0
Dir	pollux/verilog/lvs/mod10	BOM 4.0
1.3	Makefile	
1.1	mod10wrap.V	
Dir	pollux/verilog/lvs/mod11	BOM 5.0
1.4	Makefile	
1.2	mod11wrap.V	
Dir	pollux/verilog/lvs/mod12	BOM 5.0
1.4	Makefile	
1.2	mod12wrap.V	
Dir	pollux/verilog/lvs/mod3	BOM 3.0
1.2	Makefile	
1.1	mod3wrap.V	
Dir	pollux/verilog/lvs/mod8	BOM 4.0
1.3	Makefile	
1.2	mod8wrap.V	
Dir	pollux/verilog/lvs/mod9	BOM 5.0
1.3	Makefile	
1.1	mod9wrap.V	
Dir	pollux/verilog/src	BOM 7.0
1.8	.checkoutrc	
1.36	Makefile	(1.34)
1.16	Makefile.lvs	(1.14)
6.1	Makefile.newdefs	(No)
1.7	export	(1.5)
3.1	fixvref.awk	
1.9	pollux.V	
1.2	pollux.config	(1.1)
6.2	pollux.ercf	(No)
1.1	pollux.rcf	
6.1	verisim.h	
Dir	pollux/verilog/src/mod1	BOM 6.0
1.3	.checkoutrc	
1.19	Makefile	(1.17)
1.6	mod1.V	(1.5)
1.3	mod1.pif	
1.3	mod1.rcf	
5.1	module.ht	(No)
Dir	pollux/verilog/src/mod10	BOM 7.0
1.2	.checkoutrc	

1.21	Makefile	(1.19)
1.4	fp_ecl64x24.tpl	
1.7	mod10.V	
1.5	mod10.rcf	
6.1	mod10_driver.V	(No)
6.1	module.ht	(No)
Dir	pollux/verilog/src/mod11	BOM 6.0
1.2	.checkoutrc	
1.22	Makefile	(1.20)
1.11	mod11.V	
1.5	mod11.rcf	
1.3	mod11_driver.V	
1.1	mod11_ts1.V	
1.1	mod11_ts2048.V	
1.2	mod11_ts512.V	
1.2	mod11_ts64.V	
1.2	mod11_ts8.V	
1.6	mod11place.c	
5.1	module.ht	(No)
Dir	pollux/verilog/src/mod12	BOM 9.0
1.2	.checkoutrc	
1.21	Makefile	(1.18)
1.13	mod12.V	
1.3	mod12.pif	
1.4	mod12.rcf	
1.5	mod12_driver.V	(1.3)
8.1	module.ht	(No)
Dir	pollux/verilog/src/mod2	BOM 6.0
1.6	.checkoutrc	
1.18	Makefile	(1.16)
1.4	mod2.V	
1.3	mod2.pif	
1.3	mod2.rcf	
5.1	module.ht	(No)
Dir	pollux/verilog/src/mod3	BOM 6.0
1.1	.checkoutrc	
1.13	Makefile	(1.12)
1.7	mod3.V	
1.2	mod3.pif	
1.3	mod3.rcf	
5.1	module.ht	(No)
Dir	pollux/verilog/src/mod4	BOM 6.0
1.1	.checkoutrc	
1.12	Makefile	(1.10)
1.5	mod4.V	
1.1	mod4.pif	
1.3	mod4.rcf	
5.1	module.ht	(No)
Dir	pollux/verilog/src/mod5	BOM 6.0
1.1	.checkoutrc	
1.13	Makefile	(1.12)
1.5	mod5.V	
1.1	mod5.pif	
1.4	mod5.rcf	
5.1	module.ht	(No)
Dir	pollux/verilog/src/mod6	BOM 6.0
1.1	.checkoutrc	
1.9	Makefile	(1.8)
1.2	mod6.V	
1.1	mod6.pif	

```

1.2      mod6.rcf
1.1      mod6base.ly
5.1      module.ht                                (No)

Dir      pollux/verilog/src/mod7                  BOM 4.0
3.1      module.ht                                (No)

Dir      pollux/verilog/src/mod8                  BOM 6.0
1.1      .checkoutrc
1.12     Makefile                                (1.10)
1.9      mod8.V
1.2      mod8.pif
1.4      mod8.rcf
5.1      module.ht                                (No)

Dir      pollux/verilog/src/mod9                  BOM 6.0
1.2      .checkoutrc
1.16     Makefile                                (1.15)
1.5      mod9.V
1.4      mod9.rcf
1.3      mod9_driver.V                            (1.2)
1.1      mod9_ts196608.V
1.1      mod9_ts32768.V
1.1      mod9_ts4096.V
1.1      mod9_ts512.V
1.1      mod9_ts64.V
1.2      mod9_ts8.V
1.1      mod9_tsbuf.V
1.4      mod9place.c
5.1      module.ht                                (No)

====> running pollux/.checkoutrc (Fri May 12 13:00:10 PDT 1995) <====
gmake GARDS_HOST=gamorra DISPLAY=ambiorix:0.0 -C dcell dcells
gmake[1]: Entering directory `/N/rama/s9/chip-pollux/dcell'
gmake[1]: Nothing to be done for `dcells'.
gmake[1]: Leaving directory `/N/rama/s9/chip-pollux/dcell'
gmake GARDS_HOST=gamorra DISPLAY=ambiorix:0.0 -C baseplate baseplate
gmake[1]: Entering directory `/N/rama/s9/chip-pollux/baseplate'
rm -f *.ly
cd ../compass/baseplate;\
rm -f *.ly
touch clean
gmake GARDS_HOST=gamorra DISPLAY=ambiorix:0.0 floorplanmod1.sgen floorplanmod12.sgen
floorplanmod2.sgen floorplanmod3.sgen floorplanmod4.sgen floorplanmod5.sgen
floorplanmod6.sgen floorplanmod7.sgen floorplanmod8.sgen floorplanmod9.sgen
floorplanmod10.sgen floorplanmod11.sgen ../compass/baseplate/floorplanmod1.ly
../compass/baseplate/floorplanmod12.ly ../compass/baseplate/floorplanmod2.ly
../compass/baseplate/floorplanmod3.ly ../compass/baseplate/floorplanmod4.ly
../compass/baseplate/floorplanmod5.ly ../compass/baseplate/floorplanmod6.ly
../compass/baseplate/floorplanmod7.ly ../compass/baseplate/floorplanmod8.ly
../compass/baseplate/floorplanmod9.ly ../compass/baseplate/floorplanmod10.ly
../compass/baseplate/floorplanmod11.ly ../compass/baseplate/floorplanmod1_obstruct.ly
../compass/baseplate/floorplanmod12_obstruct.ly ../compass/baseplate/floorplanmod2
_obstruct.ly ../compass/baseplate/floorplanmod3_obstruct.ly
../compass/baseplate/floorplanmod4_obstruct.ly ../compass/baseplate/floorplanmod5
_obstruct.ly ../
gmake[2]: Entering directory `/N/rama/s9/chip-pollux/baseplate'
gmake[2]: Nothing to be done for `floorplanmod1.sgen'.
gmake[2]: Nothing to be done for `floorplanmod12.sgen'.
gmake[2]: Nothing to be done for `floorplanmod2.sgen'.
gmake[2]: Nothing to be done for `floorplanmod3.sgen'.
gmake[2]: Nothing to be done for `floorplanmod4.sgen'.
gmake[2]: Nothing to be done for `floorplanmod5.sgen'.
gmake[2]: Nothing to be done for `floorplanmod6.sgen'.
gmake[2]: Nothing to be done for `floorplanmod7.sgen'.
gmake[2]: Nothing to be done for `floorplanmod8.sgen'.
gmake[2]: Nothing to be done for `floorplanmod9.sgen'.
gmake[2]: Nothing to be done for `floorplanmod10.sgen'.

```

```

gmake[2]: Nothing to be done for `floorplanmod11.sgen'.
sofagen -v ../compass/vlsi.booc-dcell < floorplanmod1.sgen
mv floorplanmod1*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < floorplanmod12.sgen
mv floorplanmod12*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < floorplanmod2.sgen
mv floorplanmod2*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < floorplanmod3.sgen
mv floorplanmod3*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < floorplanmod4.sgen
mv floorplanmod4*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < floorplanmod5.sgen
mv floorplanmod5*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < floorplanmod6.sgen
mv floorplanmod6*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < floorplanmod7.sgen
mv floorplanmod7*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < floorplanmod8.sgen
mv floorplanmod8*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < floorplanmod9.sgen
mv floorplanmod9*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < floorplanmod10.sgen
mv floorplanmod10*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < floorplanmod11.sgen
mv floorplanmod11*.ly ../compass/baseplate
echo 'POBS = flatten(POBS); CPOB = flatten(CPOB); delete subcell();' \
| vlsimm -f- -v ../compass/vlsi.booc-dcell floorplanmod1 \
| sed '1s/floorplanmod1/floorplanmod1_obstruct/' > ../compass/baseplate/floorplanmod1
_obstruct.ly
echo 'POBS = flatten(POBS); CPOB = flatten(CPOB); delete subcell();' \
| vlsimm -f- -v ../compass/vlsi.booc-dcell floorplanmod12 \
| sed '1s/floorplanmod12/floorplanmod12_obstruct/' >
../compass/baseplate/floorplanmod12_obstruct.ly
echo 'POBS = flatten(POBS); CPOB = flatten(CPOB); delete subcell();' \
| vlsimm -f- -v ../compass/vlsi.booc-dcell floorplanmod2 \
| sed '1s/floorplanmod2/floorplanmod2_obstruct/' > ../compass/baseplate/floorplanmod2
_obstruct.ly
echo 'POBS = flatten(POBS); CPOB = flatten(CPOB); delete subcell();' \
| vlsimm -f- -v ../compass/vlsi.booc-dcell floorplanmod3 \
| sed '1s/floorplanmod3/floorplanmod3_obstruct/' > ../compass/baseplate/floorplanmod3
_obstruct.ly
echo 'POBS = flatten(POBS); CPOB = flatten(CPOB); delete subcell();' \
| vlsimm -f- -v ../compass/vlsi.booc-dcell floorplanmod4 \
| sed '1s/floorplanmod4/floorplanmod4_obstruct/' > ../compass/baseplate/floorplanmod4
_obstruct.ly
echo 'POBS = flatten(POBS); CPOB = flatten(CPOB); delete subcell();' \
| vlsimm -f- -v ../compass/vlsi.booc-dcell floorplanmod5 \
| sed '1s/floorplanmod5/floorplanmod5_obstruct/' > ../compass/baseplate/floorplanmod5
_obstruct.ly
echo 'POBS = flatten(POBS); CPOB = flatten(CPOB); delete subcell();' \
| vlsimm -f- -v ../compass/vlsi.booc-dcell floorplanmod6 \
| sed '1s/floorplanmod6/floorplanmod6_obstruct/' > ../compass/baseplate/floorplanmod6
_obstruct.ly
echo 'POBS = flatten(POBS); CPOB = flatten(CPOB); delete subcell();' \
| vlsimm -f- -v ../compass/vlsi.booc-dcell floorplanmod7 \
| sed '1s/floorplanmod7/floorplanmod7_obstruct/' > ../compass/baseplate/floorplanmod7
_obstruct.ly
echo 'POBS = flatten(POBS); CPOB = flatten(CPOB); delete subcell();' \
| vlsimm -f- -v ../compass/vlsi.booc-dcell floorplanmod8 \
| sed '1s/floorplanmod8/floorplanmod8_obstruct/' > ../compass/baseplate/floorplanmod8
_obstruct.ly
echo 'POBS = flatten(POBS); CPOB = flatten(CPOB); delete subcell();' \
| vlsimm -f- -v ../compass/vlsi.booc-dcell floorplanmod9 \
| sed '1s/floorplanmod9/floorplanmod9_obstruct/' > ../compass/baseplate/floorplanmod9
_obstruct.ly
echo 'POBS = flatten(POBS); CPOB = flatten(CPOB); delete subcell();' \
| vlsimm -f- -v ../compass/vlsi.booc-dcell floorplanmod10 \

```

```

| sed '1s/floorplanmod10/floorplanmod10_obstruct/' >
../compass/baseplate/floorplanmod10_obstruct.ly
echo 'POBS = flatten(POBS); CPOB = Flatten(CPOB); delete subcell();' \
| vlsimm -f- -v ../compass/vlsi.booc-dcell floorplanmod11 \
| sed '1s/floorplanmod11/floorplanmod11_obstruct/' >
../compass/baseplate/floorplanmod11_obstruct.ly
gmake [2]: Leaving directory `/N/rama/s9/chip-pollux/baseplate'
gmake GARDS_HOST=gamorra DISPLAY=ambiorix:0.0 testmod1.sgen testmod2.sgen testmod3.sgen
testmod4.sgen testmod5.sgen testmod6.sgen testmod7.sgen testmod8.sgen testmod9.sgen
testmod10.sgen testmod11.sgen testmod12.sgen testmod12.sgen baseplate.sgen floorplan.sgen
../compass/baseplate/testmod1.ly ../compass/baseplate/testmod2.ly
../compass/baseplate/testmod3.ly ../compass/baseplate/testmod4.ly
../compass/baseplate/testmod5.ly ../compass/baseplate/testmod6.ly
../compass/baseplate/testmod7.ly ../compass/baseplate/testmod8.ly
../compass/baseplate/testmod9.ly ../compass/baseplate/testmod10.ly
../compass/baseplate/testmod11.ly ../compass/baseplate/testmod12.ly
../compass/baseplate/testmod12.ly ../compass/baseplate/baseplate.ly
../compass/baseplate/floorplan.ly
gmake [2]: Entering directory `/N/rama/s9/chip-pollux/baseplate'
gawk -f pad.awk mod1.list > padlistmod1.m4
/n/rama/s9/chip-pollux/tools/bin/m4 testmod1.sgen.m4 > testmod1.sgen
gawk -f pad.awk mod2.list > padlistmod2.m4
/n/rama/s9/chip-pollux/tools/bin/m4 testmod2.sgen.m4 > testmod2.sgen
/n/rama/s9/chip-pollux/tools/bin/m4 testmod3.sgen.m4 > testmod3.sgen
/n/rama/s9/chip-pollux/tools/bin/m4 testmod4.sgen.m4 > testmod4.sgen
/n/rama/s9/chip-pollux/tools/bin/m4 testmod5.sgen.m4 > testmod5.sgen
gmake [2]: `testmod6.sgen' is up to date.
gmake [2]: `testmod7.sgen' is up to date.
gmake [2]: `testmod8.sgen' is up to date.
gmake [2]: `testmod9.sgen' is up to date.
gmake [2]: `testmod10.sgen' is up to date.
gmake [2]: `testmod11.sgen' is up to date.
gmake [2]: `testmod12.sgen' is up to date.
gmake [2]: `testmod12.sgen' is up to date.
/n/rama/s9/chip-pollux/tools/bin/m4 baseplate.sgen.m4 > baseplate.sgen
gmake [2]: `floorplan.sgen' is up to date.
sofagen -v ../compass/vlsi.booc-dcell < testmod1.sgen
mv testmod1*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < testmod2.sgen
mv testmod2*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < testmod3.sgen
mv testmod3*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < testmod4.sgen
mv testmod4*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < testmod5.sgen
mv testmod5*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < testmod6.sgen
mv testmod6*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < testmod7.sgen
mv testmod7*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < testmod8.sgen
mv testmod8*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < testmod9.sgen
mv testmod9*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < testmod10.sgen
mv testmod10*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < testmod11.sgen
mv testmod11*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < testmod12.sgen
mv testmod12*.ly ../compass/baseplate
gmake [2]: `../compass/baseplate/testmod12.ly' is up to date.
sofagen -v ../compass/vlsi.booc-dcell < baseplate.sgen
*WARNING* - Cell "floorplan" not found
mv baseplate*.ly ../compass/baseplate
sofagen -v ../compass/vlsi.booc-dcell < floorplan.sgen
mv floorplan*.ly ../compass/baseplate
gmake [2]: Leaving directory `/N/rama/s9/chip-pollux/baseplate'

```



```

gmake GARDS_HOST=gamorra DISPLAY=ambiorix:0.0 ../compass/baseplate/toplabel.ly
../compass/baseplate/toplabel.lvs.ly
gmake[2]: Entering directory `/N/rama/s9/chip-pollux/baseplate'
rm -f ../compass/baseplate/toplabel.ly
for num in 1 2 3 4 5 8 9 10 11 12 ; do \
    echo "Staring mod$num";\
    rm -f ../compass/temp.sel;\
    echo "testmod${num}_arrays" >> ../compass/temp.sel;\
    cd ../compass;\
    bubble text baseplate temp.sel ../compass/vlsi.booc-dcell | \
    nawk '{gsub(/PAD/, "PADM${num}") ; print }' | \
    nawk '{if($1 == "N" || $1 == "L") print $0;}'>baseplate/toplabel.ly;\

```

```

done
Staring mod1
Building cell tree...
.....
Cells read:      130
Instances:      2
Exploding...
Writing .ly file...
Staring mod2
Building cell tree...
.....
Cells read:      130
Instances:      2
Exploding...
Writing .ly file...
Staring mod3
Building cell tree...
.....
Cells read:      130
Instances:      2
Exploding...
Writing .ly file...
Staring mod4
Building cell tree...
.....
Cells read:      130
Instances:      2
Exploding...
Writing .ly file...
Staring mod5
Building cell tree...
.....
Cells read:      130
Instances:      2
Exploding...
Writing .ly file...
Staring mod8
Building cell tree...
.....
Cells read:      130
Instances:      2
Exploding...
Writing .ly file...
Staring mod9
Building cell tree...
.....
Cells read:      130
Instances:      2

```

```

Exploding...
Writing .ly file...
Staring mod10
Building cell tree...
.....
Cells read:      130
Instances:       2
Exploding...
Writing .ly file...
Staring mod11
Building cell tree...
.....
Cells read:      130
Instances:       2
Exploding...
Writing .ly file...
Staring mod12
Building cell tree...
.....
Cells read:      130
Instances:       2
Exploding...
Writing .ly file...
rm -f ../compass/temp.sel;
cp ../compass/baseplate/toplabel.ly ../compass/baseplate/toplabel.lvs.ly;
rm -f ../compass/temp.sel;
echo 'testmod6 arrays' >> ../compass/temp.sel;
cd ../compass;\
bubble text baseplate temp.sel ../compass/vlsi.boocell | \
nawk '{gsub(/PAD/, "PADM6") ; print }' | \
nawk '{ if($1 == "N" || $1 == "L") print $0; }' >> baseplate/toplabel.lvs.ly;
Building cell tree...
.....
Cells read:      130
Instances:       2
Exploding...
Writing .ly file...
rm -f ../compass/temp.sel;
gmake[2]: Leaving directory `/N/rama/s9/chip-pollux/baseplate'
gmake GARDS_HOST=gamorra DISPLAY=ambiorix:0.0 ../compass/baseplate/knobmap.ly
gmake[2]: Entering directory `/N/rama/s9/chip-pollux/baseplate'
rm -f ../compass/baseplate/knobmap.ly
cp knobmap.ly.orig ../compass/baseplate/knobmap.ly
gmake[2]: Leaving directory `/N/rama/s9/chip-pollux/baseplate'
gmake GARDS_HOST=gamorra DISPLAY=ambiorix:0.0 ../compass/baseplate/pollux.ly
gmake[2]: Entering directory `/N/rama/s9/chip-pollux/baseplate'
nawk '{if ($6 ~/testmod/ && $6 !~/testmod7/) { \
l=length($6);l=l-6;u=substr($6,6,l); \
print $1,$2,"0","0",$5,"\"u\"","\"\",$7,$8 ; } else { print $0 } \
}' ../compass/baseplate/baseplate.ly | sed -e 's/baseplate/pollux/g' > tmp.ly
nawk '{if ($1 != "E") print $0 }' tmp.ly > tmp2.ly;
cat ../compass/baseplate/toplabel.lvs.ly >> tmp2.ly;
echo 'E' >> tmp2.ly;
mv tmp2.ly ../compass/baseplate/pollux.ly;
rm tmp.ly;
gmake[2]: Leaving directory `/N/rama/s9/chip-pollux/baseplate'
gmake[1]: Leaving directory `/N/rama/s9/chip-pollux/baseplate'
gmake GARDS_HOST=gamorra DISPLAY=ambiorix:0.0 -C gards gards
gmake[1]: Entering directory `/N/rama/s9/chip-pollux/gards'
[ -d sofa ] || mkdir -p sofa/
cat /n/rama/s9/chip-pollux/proteus/gards/basegen/sofa_model.ly.template /n/rama/s9/chip-
pollux/compass/baseplate/toplabel.ly >sofa/sofa_model.ly
echo "E" >>sofa/sofa_model.ly

```

```

[ -d sofa/ ] || mkdir -p sofa/
/n/rama/s9/chip-pollux/proteus/gards/basegen/sofa_exclude baseplate /n/rama/s9/chip-
pollux/compass/vlsi.booc-dcell
Creating work area: /N/rama/s9/chip-pollux/gards/sofa_exclude_950512130523
Gutting mosatom_site
Gutting mobieclium_site
Gutting baseplate_clock_spars
Building sofa_exclude.ly
Removing /N/rama/s9/chip-pollux/gards/sofa_exclude_950512130523
cp sofa_exclude.ly /n/rama/s9/chip-pollux/Compass/baseplate
mv sofa_exclude.ly sofa/sofa_exclude.ly
echo './sofa/sofa_model.cdl.abgen ./sofa/sofa.pdl: \' > Depend-cdl
/n/rama/s9/chip-pollux/tools/bin/vlsimm -M -p ./sofa -v /n/rama/s9/chip-
pollux/compass/vlsi.booc-dcell sofa_model >> Depend-cdl
echo '' >> Depend-cdl
gmake[1]: Leaving directory `~N/rama/s9/chip-pollux/gards'
gmake[1]: Entering directory `~N/rama/s9/chip-pollux/gards'
[ -d sofa/ ] || mkdir -p sofa/
(echo 'CELL: ElX1;'; echo 'CELL: M1X1;'; \
  (cd /n/rama/s9/chip-pollux/proteus/gards/leaf; cat *.pdl); \
  (cd /n/rama/s9/chip-pollux/proteus/gards/sofa; cat *.pdl)) \
  grep -i 'CELL.*:[0-9]X[0-9]' \
  sed -e 's/.*[: ]\([.*[0-9][0-9]*[Xx][0-9][0-9]*\);.*[0-9]X[0-9]' \
  sort -tX +0n +1n | uniq > ./sofa/profiles.txt
sort -u /n/rama/s9/chip-pollux/dcell/list /n/rama/s9/chip-pollux/proteus/dcell/list
/n/rama/s9/chip-pollux/proteus/dcell/custom-list \
  | sed 's$/ L I (METAL2, METAL3, METAL4)/' > leaf.sel
[ -d sofa/ ] || mkdir -p sofa/
PATH=./u/chip/tools/bin:/usr/local/bin:/usr/ucb:/usr/bin:/bin:/n/rama/s9/chip-
pollux/tools/s1/bin HOME=/n/rama/s9/chip-pollux/tools \
  /n/rama/s9/chip-pollux/proteus/gards/basegen/freepos_cdl ./sofa/sofa_model.ly leaf.sel
/n/rama/s9/chip-pollux/compass/vlsi.booc-dcell
Fri May 12 13:19:48 PDT 1995
Initializing work area: /N/rama/s9/chip-pollux/gards/freepos_tmp
*****
Abstracting free cell positions with abgen
*****
Building cell tree...
Cells read: 133
Instances: 1762624
*WARNING* - Selected cell "L" not in tree
*WARNING* - Selected cell "bellybutt" not in tree
*WARNING* - Selected cell "bg3stack" not in tree
*WARNING* - Selected cell "bgbbcostm" not in tree
*WARNING* - Selected cell "bgiref" not in tree
*WARNING* - Selected cell "cache" not in tree
*WARNING* - Selected cell "cahalf" not in tree
*WARNING* - Selected cell "cerpokgen" not in tree
*WARNING* - Selected cell "cerpokgen3" not in tree
*WARNING* - Selected cell "cli" not in tree
*WARNING* - Selected cell "clkknob" not in tree
*WARNING* - Selected cell "clrepeat" not in tree
*WARNING* - Selected cell "cr" not in tree
*WARNING* - Selected cell "ctag" not in tree
*WARNING* - Selected cell "iobyte" not in tree
*WARNING* - Selected cell "iobytem" not in tree
*WARNING* - Selected cell "leddigdrv" not in tree
*WARNING* - Selected cell "ledkbip" not in tree
*WARNING* - Selected cell "ledsegdrv" not in tree
*WARNING* - Selected cell "mb" not in tree
*WARNING* - Selected cell "pl_euh" not in tree
*WARNING* - Selected cell "pl_eus" not in tree
*WARNING* - Selected cell "pl_mne" not in tree
*WARNING* - Selected cell "ttl3vnew" not in tree
*WARNING* - Selected cell "ttle2teu" not in tree
*WARNING* - Selected cell "ttle2tmn" not in tree
*WARNING* - Selected cell "ttle2ttl" not in tree

```

Exploding...

WARNING - the following cells contain metal geometry but have not been abstracted!

```
sofa_pins
pollux_ring
noisediff
ep11top
rdactop
mobieclium_noxistors
padtest
padmobi
padtestvss
padtestvdda
mobieclium_vsse
rxtesttop
padprb4m
hemmos_bs
hemmos_lrs
hemmos_ll
testmodcl
hemecl_lrs
ifr
hemifr_ts
ifcl
hemifcl_ts
mv sofa_model.cdl.abgen ./sofa/sofa_model.cdl.abgen
echo ' ' >>./sofa/sofa_model.cdl.abgen
echo '(* Cell: CGCLOCKBIAS *)' >>./sofa/sofa_model.cdl.abgen
echo 'CELLNAME: CGCLOCKBIAS;' >>./sofa/sofa_model.cdl.abgen
echo 'SIZE: 24000.24000;' >>./sofa/sofa_model.cdl.abgen
echo 'OFFSET: 0.0;' >>./sofa/sofa_model.cdl.abgen
echo 'POS: 0.0+0;' >>./sofa/sofa_model.cdl.abgen
echo 'ENDCELL;' >>./sofa/sofa_model.cdl.abgen
sed -e 's/4772000\.264000;/4772000.192000;/' \
    -e 's/3180000\.3576000;/912000.3576000;/' \
    ./sofa/sofa_model.cdl.abgen > ./sofa/sofa_model.cdl.abgen.tmp
mv ./sofa/sofa_model.cdl.abgen.tmp ./sofa/sofa_model.cdl.abgen
[ -d sofa/ ] || mkdir -p sofa/
/n/rama/s9/chip-pollux/proteus/gards/basegen/derive_bounds ./sofa/sofa_model.ly
/n/rama/s9/chip-pollux/compass/vlsi.boocell > ./sofa/sofa_bounds.gsub
/n/rama/s9/chip-pollux/proteus/gards/basegen/gsub ./sofa/sofa_bounds.gsub
</n/rama/s9/chip-pollux/proteus/gards/basegen/sofacdl.c.template \
> ./sofa/sofacdl.c
cc -g ./sofa/sofacdl.c -o ./sofa/sofacdl
./sofa/sofacdl ./sofa/profiles.txt ./sofa/sofa_exclude.ly /n/rama/s9/chip-
pollux/proteus/gards/sofa | \
sed -e '/^END_OF_FILE;/d' > sofa/sofa.cdl
cat ./sofa/sofa_model.cdl.abgen >> sofa/sofa.cdl
/usr/sbin/echo '\nEND_OF_FILE;' >> sofa/sofa.cdl
[ -d sofa/ ] || mkdir -p sofa/
PATH=.: /u/chip/tools/bin:/usr/local/bin:/usr/ucb:/usr/bin:/bin:/n/rama/s9/chip-
pollux/tools/sl/bin HOME=/n/rama/s9/chip-pollux/tools \
/n/rama/s9/chip-pollux/proteus/gards/basegen/basegen \
./sofa/sofa_model.ly /n/rama/s9/chip-pollux/proteus/gards/basegen/flat.sel
/n/rama/s9/chip-pollux/compass/vlsi.boocell
Fri May 12 13:28:50 PDT 1995
Initializing work area: /N/rama/s9/chip-pollux/gards/basegen_tmp
*****
Running abgen...
*****
Building cell tree...
Cells read: 133
Instances: 1762624
*WARNING* - Selected cell "hemmos_plug1" not in tree
*WARNING* - Selected cell "hemmos_plugu" not in tree
```

```

*WARNING* - Selected cell "hemmos_ts" not in tree
*WARNING* - Selected cell "hemmos_ul" not in tree
*WARNING* - Selected cell "padtestvddb" not in tree
*WARNING* - Selected cell "padtestvddleft" not in tree
*WARNING* - Selected cell "padtestvddright" not in tree
*WARNING* - Selected cell "padtestvddtop" not in tree
*WARNING* - Selected cell "padtestvssbot" not in tree
*WARNING* - Selected cell "padtestvssleft" not in tree
*WARNING* - Selected cell "padtestvssright" not in tree
*WARNING* - Selected cell "padtestvssstop" not in tree
*WARNING* - Selected cell "powerwaffle_e" not in tree
*WARNING* - Selected cell "hepharn" not in tree
*WARNING* - Selected cell "horm4con" not in tree
*WARNING* - Selected cell "verconm4" not in tree
*WARNING* - Selected cell "horlink" not in tree
*WARNING* - Selected cell "padhib" not in tree
*WARNING* - Selected cell "padfatwire" not in tree
*WARNING* - Selected cell "padvdd" not in tree
*WARNING* - Selected cell "padvss" not in tree
*WARNING* - Selected cell "padvdda" not in tree
*WARNING* - Selected cell "padrfr" not in tree
*WARNING* - Selected cell "padttl" not in tree
*WARNING* - Selected cell "padbl" not in tree
*WARNING* - Selected cell "padbr" not in tree
*WARNING* - Selected cell "padtl" not in tree
*WARNING* - Selected cell "padtr" not in tree
*WARNING* - Selected cell "vrrvbus" not in tree
*WARNING* - Selected cell "avddpad_m5" not in tree
*WARNING* - Selected cell "avsspad_m5" not in tree
*WARNING* - Selected cell "pllwl" not in tree
*WARNING* - Selected cell "pllwl2" not in tree
*WARNING* - Selected cell "som_m3" not in tree
*WARNING* - Selected cell "som_m4" not in tree

```

Exploding...

```

*WARNING* - the following cells contain metal
              geometry but have not been abstracted!

```

```

pollux_ring
tsensa
bgproca2d
addnf
addnfds
bgknobgen
eplltop
rdactop
ratop
aftop
rxtop
gtlb
testbb
padtestvss
mobieclium_vsse
rxtesttop
ifr
hemifr_ts
hemifcl ts

```

Abstracting flat cells for chip base...

```

Cell: hemecl_lrs
Cell: hemmos_bs
Cell: hemmos_ll
Cell: hemmos_lrs
Cell: ifcl
Cell: mobieclium_noxistors
Cell: mobieclium_probe

```

```

Cell: padmobi
Cell: padtest
Cell: padtestvdda
Cell: sofa_pins
Cell: testmodel
Cell: noisediff
*****
Compiling base cell: sofa_model
*****
Protecting targets
Simplifying metal layers
Compiling target contexts
Translation of /usr/tmp/piddles9678/sofa_model_targets.cif succeeded.
Root symbol is called ROOTCELL.
GARDS GDSGDF 7.108 -- GDS to GDF conversion
Copyright (c) 1995 SILVAR-LISCO. All rights reserved.
Design: piddles Started at: 95/05/12 13:37:11

```

```

Only specified layers will be selected
database: sofa_model_targets.gdf will be overwritten
layer file read
UOM = 1000
METRIC UNITS
** WARNING ** GDS file : last block length = 738

```

```

Terminated at      : 95/05/12 13:37:13
Elapsed CPU time   : 0 Hrs 0 Mins 0 Secs
Elapsed wall clock time : 0 Hrs 0 Mins 2 Secs
GARDS GDFPDL 7.123 -- Create PDL
Copyright (c) 1995 SILVAR-LISCO. All rights reserved.
Design: piddles Started at: 95/05/12 13:37:15

```

GDFPDL Version 7.1.23 of January 27, 1994

```

Initializing ...
Processing layout data ...
Reading name list ...
Start processing physical types ...
Writing logical to physical mapping ...
[SOFA MODEL TARGETS]
%% WARNING: Globalnet PHI_A2P not found.
%% WARNING: Globalnet PHI_B2P not found.
%% WARNING: Globalnet VREF_0PH not found.

```

```

One physical type defined.
%% WARNING: 19 Zero Length Segments.

```

```

Terminated at      : 95/05/12 13:37:20
Elapsed CPU time   : 0 Hrs 0 Mins 4 Secs
Elapsed wall clock time : 0 Hrs 0 Mins 5 Secs

```

```

Compiling routing obstructions
*****
/N/rama/s9/chip-pollux/gards/basegen_tmp
 93 -rw-r--r-- 1 chip 94829 May 12 13:37 sofa_model_base.pdl
 0 -rw-r--r-- 1 chip 0 May 12 13:37 sofa_model_leafcells.cdl
*****
Fri May 12 13:37:30 PDT 1995
sed 's/SOFA_MODEL/SOFA/' basegen_tmp/sofa_model_base.pdl > sofa/sofa.pdl
if [ -f /n/rama/s9/chip-pollux/compass/baseplate/baseplate_clock_spars.ly ] ; \
then \
  cd ./sofa; \
  /n/rama/s9/chip-pollux/proteus/gards/basegen/sofa_captiles baseplate_clock_spars
/n/rama/s9/chip-pollux/compass/vlsi.bo0-dcell ; \
fi
[ -d sofa/ ] || mkdir -p sofa/

```

```

/n/rama/s9/chip-pollux/proteus/gards/basegen/gsub \
./sofa/sofa_bounds.gsub </n/rama/s9/chip-pollux/proteus/gards/basegen/sofa.tdl.template
>sofa/sofa.tdl
gmake[1]: Leaving directory `/N/rama/s9/chip-pollux/gards'
gmake GARDS_HOST=gamorra DISPLAY=ambiorix:0.0 -C verilog pollux
gmake[1]: Entering directory `/N/rama/s9/chip-pollux/verilog'
gmake GARDS_HOST=gamorra DISPLAY=ambiorix:0.0 DISPLAY=ambiorix:0.0 -C src polluxgards
gmake[2]: Entering directory `/N/rama/s9/chip-pollux/verilog/src'
for i in mod1 mod2 mod3 mod4 mod5 mod8 mod9 mod11 mod12; do \
  gmake GARDS_HOST=gamorra DISPLAY=ambiorix:0.0 DISPLAY=ambiorix:0.0 -C ${i} vfiles ||
exit; \
done
gmake[3]: Entering directory `/N/rama/s9/chip-pollux/verilog/src/mod1'
cat /n/rama/s9/chip-pollux/proteus/verilog/diff.h mod1.v | /lib/cpp -P -C -B | sed -e '/'^
$/d'> mod1.v.tmp
mv mod1.v.tmp mod1.v
echo mod1.v | tr ' ' '\012' > vfiles
gmake[3]: Leaving directory `/N/rama/s9/chip-pollux/verilog/src/mod1'
gmake[3]: Entering directory `/N/rama/s9/chip-pollux/verilog/src/mod2'
gmake[3]: `vfiles' is up to date.
gmake[3]: Leaving directory `/N/rama/s9/chip-pollux/verilog/src/mod2'
gmake[3]: Entering directory `/N/rama/s9/chip-pollux/verilog/src/mod3'
echo mod3.v ../../../../proteus/verilog/wrappers/eplltop.v
../../../../proteus/verilog/wrappers/gedeplltop.v | tr ' ' '\012' > vfiles
gmake[3]: Leaving directory `/N/rama/s9/chip-pollux/verilog/src/mod3'
gmake[3]: Entering directory `/N/rama/s9/chip-pollux/verilog/src/mod4'
gmake[3]: `vfiles' is up to date.
gmake[3]: Leaving directory `/N/rama/s9/chip-pollux/verilog/src/mod4'
gmake[3]: Entering directory `/N/rama/s9/chip-pollux/verilog/src/mod5'
gmake[3]: `vfiles' is up to date.
gmake[3]: Leaving directory `/N/rama/s9/chip-pollux/verilog/src/mod5'
gmake[3]: Entering directory `/N/rama/s9/chip-pollux/verilog/src/mod8'
gmake[3]: `vfiles' is up to date.
gmake[3]: Leaving directory `/N/rama/s9/chip-pollux/verilog/src/mod8'
gmake[3]: Entering directory `/N/rama/s9/chip-pollux/verilog/src/mod9'
gmake[3]: `vfiles' is up to date.
gmake[3]: Leaving directory `/N/rama/s9/chip-pollux/verilog/src/mod9'
gmake[3]: Entering directory `/N/rama/s9/chip-pollux/verilog/src/mod11'
gmake[3]: `vfiles' is up to date.
gmake[3]: Leaving directory `/N/rama/s9/chip-pollux/verilog/src/mod11'
gmake[3]: Entering directory `/N/rama/s9/chip-pollux/verilog/src/mod12'
gmake[3]: `vfiles' is up to date.
gmake[3]: Leaving directory `/N/rama/s9/chip-pollux/verilog/src/mod12'
for i in mod1 mod2 mod3 mod4 mod5 mod6 mod8 mod9 mod10 mod11 mod12; do \
  gmake GARDS_HOST=gamorra DISPLAY=ambiorix:0.0 DISPLAY=ambiorix:0.0
DISPLAY=ambiorix:0.0 -C ${i} ${i}gards || exit; \
rm -f ${i}/gards/garout.cfg; \
rm -f ${i}/gards/gplace.boc; \
rm -f ${i}/gards/gplace.cfg; \
rm -f ${i}/gards/gplace.defn; \
rm -f ${i}/gards/gplace.nic; \
rm -f ${i}/gards/gplace.noc; \
rm -f ${i}/gards/maskout.cfg; \
rm -f ${i}/gards/pgroute.cfg; \
rm -f ${i}/gards/pcomp.cfg; \
rm -f ${i}/gards/slnet.boc; \
rm -f ${i}/gards/${i}.bacdly; \
rm -f ${i}/gards/${i}.ctrl; \
rm -f ${i}/gards/${i}.gil; \
rm -f ${i}/gards/${i}.ndb; \
rm -f ${i}/gards/${i}.nof; \
rm -f ${i}/gards/${i}.pif; \
rm -f ${i}/gards/${i}.pof; \
rm -f ${i}/gards/${i}.rcf; \
rm -f ${i}/gards/${i}.spn; \
rm -f ${i}/gards/${i}.vrf; \
rm -f ${i}/gards/${i}.xrf; \

```


[illegible]

[illegible]


```
cp /n/rama/s9/chip-pollux/ged/startup.concept startup.concept
echo 'use local.wrk' >> startup.concept
/n/rama/s9/chip-pollux/tools/bin/shebody -e mod1.lvsedif -V -F -p -w local.wrk -g ./ -j
mod1.kmap \
-1 /n/rama/s9/chip-pollux/proteus/spice/leaf \
-1 /n/rama/s9/chip-pollux/proteus/exlax/edif
Generating body for mod1...
Generating spice.cn1
Merging edif libraris. This will take a few minutes...
```

Running emerge compiled on Fri Apr 28 03:07:07 GMT 1995

```
rm -rf xshadowx
/n/rama/s9/chip-pollux/tools/bin/ged2lvs -O -l mod1
```

MicroUnity Spice Interface Version 1.94
(C) Copyright (C) 1990 Valid Logic Systems, Inc.

Processing Scald directories

/□

```
-□
\□
|□
ld.so: warning: /usr/lib/libc.so.1.6.1 has older revision than expected 8
```

Cadence Design Systems, Inc. ValidPAGECOMP 02.2 sun4-p02
(C) Copyright 1982,1992 Cadence Design Systems, Inc.

Reading the compiler directives.
Compiler directives read (00:00:02.48)

Reading text macro definitions.
Text macro definitions read (00:00:00.06)

Reading property attributes.
Property attributes read (00:00:00.05)

Compiling MOD1.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.78)

Compiling FLAG.PART.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.03)

Compiling TSENSA.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.25)

Compiling BGPROCA2D.SPICE.1.1
No parameters

No errors detected

No oversights detected
No warnings detected

Page compiled (00:00:01.15)

Compiling ADDNF.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.56)

Compiling ADDNFDS.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.70)

Compiling TSAD2SLOPE.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.36)

Compiling TSBANDGAP.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:01.14)

Compiling BGPLEV.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.68)

Compiling TSALARM.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.51)

Compiling BGCSTG2.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.11)

Compiling BGCOMP2.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.31)

Compiling ADDIFOA.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.43)

Compiling ADCAP2.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.73)

Compiling ADR50K.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.21)

Compiling BGPA2DBI.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.31)

Compiling XCNAND4C.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.21)

Compiling XCINVC.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.06)

Compiling XCNAND2C.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.11)

Compiling BGPREAMP.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.23)

Compiling XCNAND3C.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.16)

Compiling BGISVR.SPICE.1.1
No parameters

No errors detected
1 oversight detected
No warnings detected

Page compiled (00:00:00.56)

Compiling AD4INV.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.23)

Compiling AD1BDAC.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.80)

Compiling ADCAP1.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.33)

Compiling BGPREF.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.33)

Compiling AD1BDACN.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:01.33)

Compiling ADMUTE.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.18)

Compiling ADBIAS.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.50)

Compiling TSADCOMP.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.28)

Compiling TSADOA.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.41)

Compiling PLD.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.10)

Compiling ADR480K.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.40)

Compiling BGPROAP.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.13)

Compiling BGPROAN.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.10)

Compiling XCNC.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.05)

Compiling XCPC.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.05)

Compiling BGPSU.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.15)

Compiling BJT1.SPICE.1.1
No parameters

No errors detected
No oversights detected
1 warning detected

Page compiled (00:00:00.03)

Compiling BGRES2.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.31)

Compiling BGRES3.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.43)

Compiling XCNCPRIM.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.06)

Compiling XCPCPRIM.SPICE.1.1
No parameters

No errors detected
No oversights detected
No warnings detected

Page compiled (00:00:00.08)

Start time = 13:58:26.00
Ending time = 13:59:56.00
Elapsed time = 00:01:30.00
CPU time = 00:00:28.93

<lnk>#1 ERROR(145): Pin name does not exist
Drawing: "TSENSA".SPICE.1.1
No parameters
Body: "TSBANDGAP" (Path="5P")
Unfound pin: "VPP"

Drawing: "MOD1".SPICE.1.1
No parameters
Body: "TSENSA" (Path="37P")
Pins of the body:
"CRB0_AB0PM"
"CRB1_AB0PM"
"DUN_AB0PM"
"TSRSTN_AB0PM"
"SEL_AB0PM"
"MLTD_AB0PM"
"COMP_AB0PM"
"BGV_V"
"RAMP_V"
"ADIN_V"
"OTCI_V"

<lnk>#2 ERROR(145): Pin name does not exist
Drawing: "ADDNF".SPICE.1.1
No parameters
Body: "ADDIFOA" (Path="7P")
Unfound pin: "VPP"

Drawing: "MOD1".SPICE.1.1
No parameters
Body: "ADDNF" (Path="39P")
Pins of the body:
"DIN ABD0PF"
"DIN ABND0PF"
"CD_AB0PM"<0>
"CD_AB0PM"<1>
"CD_AB0PM"<2>
"CD_AB0PM"<3>

"REF1P5_V"
"ADOUTP_V"
"ADOUTN_V"

<lnk>#3 ERROR(145): Pin name does not exist
Drawing: "ADDNFDS".SPICE.1.1
No parameters
Body: "RDIFFS" (Path="8P")
Unfound pin: "VFP"

Drawing: "MOD1".SPICE.1.1
No parameters
Body: "ADDNFDS" (Path="40P")
Pins of the body:
"DIN_ABD0PF"
"DIN_ABND0PF"
"CD_AB0PM"<0>
"CD_AB0PM"<1>
"CD_AB0PM"<2>
"CD_AB0PM"<3>
"REF1P5_V"
"ADOUTP_V"
"ADOUTN_V"

<lnk>#4 ERROR(145): Pin name does not exist
Drawing: "ADDNFDS".SPICE.1.1
No parameters
Body: "AD4INV" (Path="60P")
Unfound pin: "CD_ABM"<0>

Drawing: "MOD1".SPICE.1.1
No parameters
Body: "ADDNFDS" (Path="40P")

<lnk>#5 ERROR(145): Pin name does not exist
Drawing: "ADDNFDS".SPICE.1.1
No parameters
Body: "ADMUTE" (Path="61P")
Unfound pin: "MUTE_ABM"

Drawing: "MOD1".SPICE.1.1
No parameters
Body: "ADDNFDS" (Path="40P")

<lnk>#6 ERROR(145): Pin name does not exist
Drawing: "ADDNFDS".SPICE.1.1
No parameters
Body: "AD1BDACN" (Path="55P")
Unfound pin: "ADG_ABM"<0>

Drawing: "MOD1".SPICE.1.1
No parameters
Body: "ADDNFDS" (Path="40P")

(00:00:04.93)
Calling ValidCompiler ...
(00:01:33.91)

Fatal error(s) found while compiling
Please correct the design and rerun the program

MicroUnity Spice Interface run on May 12 13:58:17 1995
DESIGN NAME : 'MOD1'
DESIGN COMPILATION ON May 12 13:58:23 1995

2 errors detected
No oversight detected
1 warnings detected

cpu time 0:00:04
elapsed time 0:01:39

ld.so: warning: /usr/lib/libc.so.1.6.1 has older revision than expected 8

Cadence Design Systems, Inc. ValidCOMPERR 02.2 sun4-p02
(C) Copyright 1982,1992 Cadence Design Systems, Inc.

Reading the compiler directives.
Compiler directives read (00:00:02.34)

Reading text macro definitions.
Text macro definitions read (00:00:00.01)

Reading property attributes.
Property attributes read (00:00:00.08)

Retrieving listing files.
Listing files retrieved (00:00:07.91)

Starting output of error documentation.
End of error help message output (00:00:00.20)

Start time = 13:59:58.00
Ending time = 14:00:28.00
Elapsed time = 00:00:30.00
CPU time = 00:00:10.59

*** /n/rama/s9/chip-pollux/tools/bin/ged2lvs failed
gmake[4]: *** [mod1.splvs] Error 1
rm mod1.kmap mod1.lvsedif
gmake[4]: Leaving directory `/N/rama/s9/chip-pollux/verilog/src/mod1'
gmake[3]: *** [modlgards] Error 1
gmake[3]: Leaving directory `/N/rama/s9/chip-pollux/verilog/src/mod1'
gmake[2]: *** [subdirs] Error 1
gmake[2]: Leaving directory `/N/rama/s9/chip-pollux/verilog/src'
gmake[1]: *** [pollux] Error 1
gmake[1]: Leaving directory `/N/rama/s9/chip-pollux/verilog'
gmake: *** [pollux] Error 1
[finished at Fri May 12 14:00:33 PDT 1995 -- exit status 1]

From: pmayer (Patricia Mayer)
Sent: Friday, May 12, 1995 12:16 PM
To: howard
Cc: tbr
Subject: re: test board schedule

FYI

I expect Naug Le to be starting on the 22nd (not confirmed). So be prepared!
-Pattie

> From warren@godzilla Fri May 12 09:42:04 1995
> Date: Fri, 12 May 1995 09:42:03 -0700
> From: warren@godzilla (Mark Warren)
> To: pmayer@godzilla
> Subject: re: test board schedule
> Content-Length: 557
>
> Pattie -
>
> Work falls into several categories -
>
> Round Probe Cards
> - euterpe tape out June 19
> - mnenosyne tape out June 19
>
> Pandora Board for Cronus
> - special one for cronus die tape out June 26
>
> Test Performance Boards
> - euterpe tape out June 12
> - mnenosyne tape out June 12
> - cronus die tape out June 19
>
> So we're talking 2 rather simple probe card pcbs, a pandora
> modification and three test performance boards.
> This is a lot of work - I would estimate about 3 man weeks at least.
>
> We need to get started very very soon!!
>
> Mark
>

From: pmayer (Patricia Mayer)
Sent: Thursday, May 11, 1995 11:41 PM
To: tbr; tbe; philip; dbulfer; woody; howard; pmayer
Cc: albers
Subject: Re: PCB layout schedule

> From pmayer Sun May 7 15:59:39 1995
> To: tbr, tbe, philip, dbulfer, woody, howard
> Subject: PCB layout schedule
> Cc: pmayer, albers
>
>> I made a mistake on the May 5th PCB layout schedule.
>
> I have Euterpe, Mnemo and Herminator due on the 15, 16 and 17. I ment
> 5, 6 and 7 but on second thought, I'd like to give each board 2 days.
> I've asked Howard to change the planes to be positive so we can review
> them without having to photoplot the boards.
>
> Perhaps we can have one design review for all three boards Friday 3:00?
>
> Please let me know if there are and scheduling conflicts.
>
> Thanks
> -Pattie

The reality is the boards really won't be ready until Wednesday due to the extensive mechanical edits. Because of Tims staff meeting on Wednesday we should postpone the review until Thursday 2:00?

Please let me know if there are any scheduling conflicts?

Thanks
-Pattie

From: hopper (Mark Hofmann)
Sent: Thursday, May 11, 1995 6:09 PM
To: tbr (Tim B. Robinson); geert (Geert Rosseel)
Cc: cadettes
Subject: More detailed estimate of CAD tasks to tapeout Euterpe

Hi,

The CAD folks met today to try to quantify the amount of work remaining between now and a fule Euterpe maskout. The numbers below assume that the DRC flow is stable and that no re-wrok will be needed. So it should be taken as a minimum guideline.

Weeks	Task
0	SDEC synthesis
2	Metal Waffle
2	Metal synthesis
0	DRC flow maintenance
1.5	Frame cell mods
1.5	scribe fill cells
1	scribe/die interface
1	copyright/dev ids
1	Eu hand route
2x6C	Eu DRC
2x6C	Po DRC
1.5C	Eu LVS
1.5C	Po LVS
1	Fix Eu LVS
2	mod. Mobi fracture flow
5-6x2C	fracture Eu
5-6x2C	fracture Po
3	Vlsimm enhancements
0	Eu/tools snap
0	Eu csyn
1	Eu celltest
0.5	Eu Space xformer
0.5	via enlargement check
3	manual SDEC fix
0	mobieclium noxisters
0.5	sofagen/make Po (+ Rich/Johnny)
8	Rev Po to current Mobi rules
0.5	fracture Eu space xformer
2	gen space xformer frame
---	-----

32 Man weeks / 4 people fulltime => 8 weeks
51 CPU weeks / 4 machines in parallel => 13 weeks

Assume that CPU cycles prior to fracture can be overlapped with man weeks for about a half overlap or 6.5 weeks.

Therefore the first mask can be created in about 8 weeks and the last mask will be finished in 8 + 6.5 = 14.5 weeks

-hopper

From: geert (Geert Rosseel)
Sent: Thursday, May 11, 1995 11:54 AM
To: fung
Cc: ahn; al; cadettes; hopper; mouss; paulp; tbr; ted
Subject: RE: Mask vendors are ready to go...

Fung says :

Let us not repeat the error we made on the device 0001. Since doing so we will NOT gain anything and instead we will lose on the delivery schedule and possibly more than double the cost of the mask set. That is, we will end up paying the price for two mask sets and with a late delivery.

What we were aiming for is a single mask set that allows us to debug our circuits at all levels : single devices, metal and contact parametrics, small memories, ringoscillators, some circuits and Euterpe.

Our inability to get anything out of the current Castor and Pollux has made it impossible for us to debug any circuits. If we only tape-out a euterpe, it will work or it won't. If it doesn't, there is currently NO WAY we can find out why. We will not even have basic parametrics such as wire and contact resistances,

If it is a lot cheaper, as you say, to tape out a separate Euterpe Mnemo, Pollux and Castor and push these 4 designs through the fab, then this might be a viable option. Especially now that we have 3 vendors ready to take tapes.

However, if we do that, Castor and Pollux should get higher priority in the mask-making and in the fab than our designs.

From: fung [fung@charybdis]
Sent: Thursday, May 11, 1995 11:25 AM
To: Mark Hofmann
Cc: Anh Ngo; cadettes; Geert Rosseel; Paul Poenisch; Tim B. Robinson; al@charybdis; mouss@charybdis; ted@charybdis
Subject: RE: Mask vendors are ready to go...

Hopper writes:

Prerequisites:

1. A completely placed and routed Euterpe and Pollux (2 Euterpe and 2 Pollux on this next tapeout are planned.) We are quite close, (but not yet done), to a fully routed Euterpe that makes timing. Once we have that we can run DRC and LVS on the full chip. These may overlap. The LVS has the longer running time. It now takes about 5 days of CPU. Assume we need two runs to make sure Euterpe is clean. Assume we can overlap the last LVS with initial mask generation (if that LVS fails, we need to restart the masks).

Hopper,

Based on our experience with our device 0001 (a combo of Calliope 0, Pollux, and Castor), we found it really made very very difficult for the mask vendor to place more than one device patterns onto the same glass. Since the the first try, our mask vendors have voiced their strong opinion to us regarding the combo type of mask patterns. They felt that the combo mask will only increase the failure rate and directly impact the mask yield.

Unlike making wafers, mask making is an effort of making the "master pattern" from the data tape. It must be defect free and with all the tight control of critical dimension and registration. To make a single master pattern and repeat it for four times on the same glass has much more chance to success than try to make two master patterns and repeat two times each. The level of difficulty is more than 2X since we now want two good master patterns on the SAME glass.

Let us not repeat the error we made on the device 0001. Since doing so we will NOT gain anything and instead we will lose on the delivery schedule and possibly more than double the cost of the mask set. That is, we will end up paying the price for two mask sets and with a late delivery.

Fung

From: hopper (Mark Hofmann)
Sent: Wednesday, May 10, 1995 4:49 PM
To: Tim B. Robinson
Subject: Re: Mask vendors are ready to go...

Tim B. Robinson writes:

Thanks. I'd certainly like to understand better in the morning. The main concern I had is that I expected your mail to come as a bombshell to mouss. At his staff meeting this morning I had felt obliged to inform him of a fairly significant slip on mnemosyne, mostly because the verification folks have been delayed switching over from euterpe. This prompted ahn to ask about euterpe, which geert described as really close, modulo further design rule changes. Neither of us mentioned or expressed great concern about the CAD task, and I certainly had no appreciation of what is being asked here.

I have a real concern that we now have something so complex that even if we can make it work it will be essentially unusable. We have to be able to compete with people who can get 30 day turns from a foundry including masks. With the present estimates we'll be lucky to get 2 turns per year and at that rate anything we do build is likely to be out of date before we can get it to volume. So, we should be seriously assessing the true value of some of the aggressive features of the process. In particular, are we making a local optimization here which is having a negative impact on our operations as a whole?

I agree with you, Tim.

Tom and Dave and I discussed briefly new ways of parallelizing the synthesis steps (by using quadrants, forking processes, etc). Nobody feels too comfortable with these approaches at present. I asked Al if this process would ever simplify. He did not think so- as soon as we get this working we would go to .4u and things would get complicated again. The problem is simply that the raw number of rectangles that require processing is huge. In fact, another concern we raised is that we are right at the addressability limit of a 32-bit machine. We're concerned that we may need more than 2gig of swap space per process for some of this work (since we have to hold multiple layers in memory at the same time now).

-hopper

From: tbr
Sent: Wednesday, May 10, 1995 11:32 PM
To: hopper (Mark Hofmann)
Subject: Re: Mask vendors are ready to go...

Mark Hofmann wrote (on Wed May 10):

Tim B. Robinson writes:

Mark, I'm too much out of the loop again on an issue of this magnitude, and I would have appreciated being brought up to speed on the latest mask generation issues before a posting as wide as this went out. Can we discuss first thing in the morning please?

It would appear to me from what you describe here that some of the requirements from the fab have now become so onerous as to be unworkable. Clearly if it is going to take 8 weeks of serial time to generate a maskset, followed by mask generation, we appear to have lost any hope of being able to debug our designs in a reasonable time. I anticipate your posting will attract mouss's attention in a big way and I need to be fully up to speed on what's going on.

Thanks
Tim

Tim,

Okay, I'm sorry if I've done things out of order here. I did look for you earlier this afternoon before posting. Sure, let's go over this tomorrow morning. I mailed my reply to the same folks that Fung addressed.

Let me try, briefly, to catch you up on some status now:

We have been having regular Wednesday meetings with Al and Paul but not Fung. We had one today. In that meeting we (Dave, Tom, Kurt, Dan and I were all in attendance) expressed our concern that the back-end synthesis effort was approaching the time it takes to get a wafer through the fab. This was discussed in that meeting, so my estimate of 8 weeks CPU (which was subsequently discussed in detail with Dave, Tom and Kurt) should not surprise Al or Paul. What has happened is that much that we could do in parallel we now have to do in serial because the design rules have become so complex as to involve many layers at once (the number of dependencies on the metal layers has gone up very much). We voiced this at the Fab meeting today.

I tried to be as straight-forward in my mail as possible. If anything, I think I have underestimated the amount of effort needed (both CPU and human) to complete this tapeout. We discussed simplifying the design rules- but Paul and Al pushed back on this saying 1. This is a complex process, if it were simple then everyone could do it. 2. They would rather push a task onto the CAD guys than the mask layout people. I agree with point 2.

The difficulty is that each week the Fab group drops another bombshell on us. This week it is a new synthesis of the iso layer (150). The problem is that we have not yet recovered from the bombshells of previous weeks. CAD can do all the tasks that the Fab group is asking of us, but we can't do it infinitely fast. We feel the ground isn't stable under us. The design rules change every week. Every week many new rules are added.

Adding CPU to trex or upgrading to a Power Challenge (from our Challenge) might help out on the CPU side.

I'm not complaining. I'm just trying to state the facts and respond to Fung's request for a date when we could provide our first masks to the mask houses.

Thanks. I'd certainly like to understand better in the morning. The main concern I had is that I expected your mail to come as a bombshell to mouss. At his staff meeting this morning I had felt obliged to inform him of a fairly significant slip on mnemosyne, mostly because the verification folks have been delayed switching over from euterpe. This prompted ahn to ask about euterpe, which geert described as really close, modulo further design rule changes. Neither of us mentioned or expressed great concern about the CAD task, and I certainly had no appreciation of what is being asked here.

I have a real concern that we now have something so complex that even if we can make it work it will be essentially unusable. We have to be able to compete with people who can get 30 day turns from a foundry including masks. With the present estimates we'll be lucky to get 2 turns per year and at that rate anything we do build is likely to be out of date before we can get it to volume. So, we should be seriously assessing the true value of some of the aggressive features of the process. In particular, are we making a local optimization here which is having a negative impact on our operations as a whole?

Tim

From: hopper (Mark Hofmann)
Sent: Wednesday, May 10, 1995 9:57 PM
To: Tim B. Robinson
Subject: Re: Mask vendors are ready to go...

Tim B. Robinson writes:

Mark, I'm too much out of the loop again on an issue of this magnitude, and I would have appreciated being brought up to speed on the latest mask generation issues before a posting as wide as this went out. Can we discuss first thing in the morning please?

It would appear to me from what you describe here that some of the requirements from the fab have now become so onerous as to be unworkable. Clearly if it is going to take 8 weeks of serial time to generate a maskset, followed by mask generation, we appear to have lost any hope of being able to debug our designs in a reasonable time. I anticipate your posting will attract mouss's attention in a big way and I need to be fully up to speed on what's going on.

Thanks
Tim

Tim,

Okay, I'm sorry if I've done things out of order here. I did look for you earlier this afternoon before posting. Sure, let's go over this tomorrow morning. I mailed my reply to the same folks that Fung addressed.

Let me try, briefly, to catch you up on some status now:

We have been having regular Wednesday meetings with Al and Paul but not Fung. We had one today. In that meeting we (Dave, Tom, Kurt, Dan and I were all in attendance) expressed our concern that the back-end synthesis effort was approaching the time it takes to get a wafer through the fab. This was discussed in that meeting, so my estimate of 8 weeks CPU (which was subsequently discussed in detail with Dave, Tom and Kurt) should not surprise Al or Paul. What has happened is that much that we could do in parallel we now have to do in serial because the design rules have become so complex as to involve many layers at once (the number of dependencies on the metal layers has gone up very much). We voiced this at the Fab meeting today.

I tried to be as straight-forward in my mail as possible. If anything, I think I have underestimated the amount of effort needed (both CPU and human) to complete this tapeout. We discussed simplifying the design rules- but Paul and Al pushed back on this saying 1. This is a complex process, if it were simple then everyone could do it. 2. They would rather push a task onto the CAD guys than the mask layout people. I agree with point 2.

The difficulty is that each week the Fab group drops another bombshell on us. This week it is a new synthesis of the iso layer (150). The problem is that we have not yet recovered from the bombshells of previous weeks. CAD can do all the tasks that the Fab group is asking of us, but we can't do it infinitely fast. We feel the ground isn't stable under us. The design rules change every week. Every week many new rules are added.

Adding CPU to relax or upgrading to a Power Challenge (from our Challenge) might help out on the CPU side.

I'm not complaining. I'm just trying to state the facts and respond to Fung's request for a date when we could provide our first masks to the mask houses.

-hopper

From: tbr
Sent: Wednesday, May 10, 1995 8:49 PM
To: pmayer (Patricia Mayer)
Cc: pmayer
Subject: PCB layout schedules

Patricia Mayer wrote (on Wed May 10):

Tim,

Just wanted to let you know whats going on.

Have one more ECO for Hestia-Main today and high hopes of finishing by Monday at our 3:00 meeting. (and with a sinus infection!)

Euterpe, Mnemo and Herminator have ECO's pending with major mechanical changes. The edits are easy enough, just time consuming, like move the 80 pin connector over just a tad... Thats about 100 edits. I foresee a design review by the middle of next week.

Howard was very receptive to doing the test boards, even offered to start the back plane and then hand it over to the next person.

I hope to move after Hestia is done, hopefully by the beginning of the week and then I need to set up the other cube for the contractor. Did we get that one? Who else should I contact for that?

Yes, though we have to share with the additional book case/file cabinet in there for the library overflow. You may want to touch base with juli to co-ordinate. If it's just a matter of moving terminals etc, ask sysadmin. If you need boxes etc armando. We also need to make sure we have an NDA with the contractor before disclosing anything and go through the basic orientation os security etc before getting started.

Tim

From: doi (Derek Iverson)
Sent: Wednesday, May 10, 1995 7:18 PM
To: guarino; gmo; doi; gregg; claseman; lisa; jeffm; iimura; deborah
Cc: hestia
Subject: Software Bringup Meeting Minutes - May 10, 1995

Software Bringup Meeting

May 10, 1995

Next Meeting: May 17 at 2:00 pm.

Attendees: guarino, gmo, doi, gregg, claseman, lisa,
jeffm, iimura, deborah

New Action Items

Item: Preparation for 'How to Debug Euterpe' presentation
Who: jeffm, doi, gregg, others?
Status: In Progress

05/10 Need to put together a packet of information about the tools
and utilities used for effective debugging.
std tool (create source listing annotated with trace
information)
mkimg (building load images and files suitable for
loading in the HW simulators)
likelevel log description
... and more.

Review of Action Items

Item: When do we have a full calliope simulation available (IKOS)?
Who: lisar
Status: Pending

04/26 This topic was raised as we talked about when the Snap/Unsnap
item should be brought back from the Suspended list.
05/10 Lisar was not able to attend the meeting.

Item: Modify startup code in stb/stand to read the cerberus node number
Who: gmo
Status: Pending

05/03 No new progress
05/10 Ditto.

Item: Plan for testing remote debugging environment
Who: everyone
Status: Pending

04/13 This will be discussed at the next meeting.
04/20 Short discussion on the pieces required and if we want to have
a standalone version of the hostio software.
Wally has started on this already.

04/26 Any work with snoopy/dram models will begin after euterpe tapeout when jeff has more time.
05/03 Wayne F. expects to have a prototype ISA card ready for testing tomorrow. Gmo is in the process of writing code to test the card.
05/10 Gmo has written a test program. Gmo and Wayne need to get together and try it on the board.

Item: ukernel needs to detect machine checks and 'do the right thing'
Who: lisa
Status: Pending

04/13 No new progress.
04/20 On list of things to do. Lower priority.
04/26 Ditto.
05/03 Ditto.
05/10 Ditto.

Item: Modify tests in diag tree to use tcc instead of tgcc
Who: guarino, jeffm, doi
Status: Pending

03/29 Loretta has changed the diag tests to use tcc instead of tgcc but has not checked in the changes. Lisar wants these changes to be made after we are able to run the tests successfully using tgcc.
04/20 Loretta has the stand/diag tests converted to use TCC. Derek is going to modify the Makefiles in stand/diag/memhi so that the programs that required tgcc will have an explicit rule.
04/26 Loretta is ready to check in the TCC related changes. Derek is to check with lisar if this is OK.
05/03 Loretta was ready to check in the changes but a last minute test uncovered a compiler bug. Checkin will happen after the tests build and run OK.

Item: Unsnap code
Who: lisar, jeffm
Status: In progress.

02/15 The issue of restarting the hardware from an IKOS dump was discussed and the need for an architectural snap/unsnap facility was questioned. Since the meeting it has been re-discovered (jeffm wasn't there to remind us of an earlier decision) that we are planning on loading architectural state into an IKOS simulation and not from a total IKOS logic dump. We also determined that when it came time to run some of the larger tests (real-time benchmark) we would need the capability to start an IKOS simulation from an architectural dump anyhow.
03/01 For the short term we are going to focus on a simpler approach for loading and running DVTs, the kernel, and kernel tests. This item will likely come back in April.
05/10 Back to life. Does the IKOS support RAM dump?

Item: Create performance test plan
Who: claseman
Status: [11/30] No progress as focus is on functionality.

11/30 We continue to run tests to help us compare terp vs hardware performance. We still need to put together the actual performance tests that need to be run on the hardware.
05/10 Back to life. Tim has checked in more tests. We are going to use

the tests written for hermes and cerberus read accesses in the 'How to debug Euterpe' presentation on monday. We will be looking for the time it takes from the instruction issue, entry to nb, and completion of access.

Suspended Items

None.

Completed Items

Item: Schedule a meeting to go over how to run stuff on HW simulator and read HW traces.

Who: jeffm, lisar, doi

Status: Done.

04/05 doi is going to talk to lisar about when this would be appropriate. Lisar has volunteered to show up at the next meeting and discuss this.

04/13 lisar presented a strategy that will give other the ability to run and do software debug with the hardware simulators. Jeffm to write up a crib sheet that will aid with debugging in the hardware accelerator environment.

04/20 Dave Tomlinson will require such a session. Sometime after May 8.

04/26 Suspended until 05/08.

05/10 Meeting scheduled for 05/15 at 1:00pm.

Terp Feature Status

inprog o Add support for host I/O through the sdram
done o Holes in address spaces => machine checks, exceptions, etc.
inprog o Ifetch protection granularity
 - performance vrs accuracy tradeoff
inprog o Fetch instructions as octlets
inprog o Accuracy wrt HW simulator(s?)
 o Better latency model for Calliope accesses
 o Implement hardware configuration through Cerberus regs
 (SDRAM parameters, dram enable?)
 o Checkpoints/Snapshots
 o hermes and cerberus timeout machine checks
 - does jeff have a test for this? Currently the TSA is
 followed (immediate mchk) instead of waiting for watchdog
 timeout.
inprog o ability for terp to load hermes sections
 - lisar would like this functionality added
 o remove stbio from hwterp.

Performance Test Status

Tim posted the following numbers for loads from dram and rom to the euterpe newsgroup.

load from dram: MINOR CYCLE COUNT : 0x00000ce

load from rom : MINOR CYCLE COUNT : 0x00002f4

Here is the current (un-prioritized) list of desired performance analysis.

Tests to be created:

- 1) stores/loads to dram, hermes, cerberus, load from rom

- 2) icachemiss
- 3) dcachemiss
- 4) load ltlb entry (write+read)
- 5) load gtlb entry (write+read)
- 6) NB overflow
- 7) generate an interrupt
- 8) return from interrupt
- 9) multiple cylinders trying to take exceptions at the same time
- 10) predicated branch
- 11) unpredicted branch
- 12) One instruction from each instruction class:

```

arith    store    eshiftx    imul64
branch   swap     epop       idiv
branchx  stored   imul4      gfmul
load     eset     imul8      bgate
loadx    eshift   imul16     atomic ops
nbload   eshiftx  imul32

```

- 13) Inner loop sequences:
 - cable in main handler-- inner loop in EQ_UPDATE_WEIGHTS() (khp)
 - IDCT code (fur)
 - pseudo-Huffman decode loop (fur)
 - a reconstruction routine from macroblock.c (fur)
 - NTSC encode loop (fur)
 - rs_compute_syndrome() (fur)

Test Status and General Discussion

The nullTest is getting further. A bug was found in the microkernel at about 20,000,000 'ticks'. Gmo fixed the problem and we started running the nullTest again.

A 'broken' version of cachenastry4 found a HW bug this morning.

There are still about 20 tests that need to be written.

From: Mark Hofmann [hopper@gaea.microunity.com]
Sent: Wednesday, May 10, 1995 7:08 PM
To: fung
Cc: geert@gaea; anh@gaea; cadettes@gaea; paulp@gaea; tbr@gaea; al@charybdis; fung@charybdis; mouss@charybdis; ted@charybdis
Subject: Re: Mask vendors are ready to go...

fung writes:
Geert,

We are very happy to inform you that our mask vendors are ready to take our reticle jobs at any time now. Under our incentivized purchasing program, they are committed to ship to us with a schedule of at least one mask per day! Currently we now have two completely qualified mask vendors with a possibility of getting the third one qualified before Q4 this year. It is our goal that by the end of this year, if we choose to, we should be able to tape out three mask sets at the same time!

Since we have worked very closely with our mask vendors, they have submitted a request to us and needed an estimated answer. That is, in order to help them plan for allocating the capacity, they would like to have an estimated tape-out dates for both the Euterpe and Mnemo. It will be very thankful if you can provide the information to us as soon as you can.

To give you a very brief background: for the last three to four months, Kurt, Al, & myself have continued to work with our two major mask suppliers for perfecting the reticle making process. According to the most recent results we have seen from the two mask vendors -- DuPont (Santa Clara) & Photronics (Sunnyvale) -- we felt that the reticle processing technology have finally become much more production worthy to suit our need. The potential supplier -- Align Rite Corp. -- has come in a long way and making a very good progress. It has narrowed in their shortcomings in terms of the processing equipment. Major acquisitions will be made by the end of this month in order to have the required capability.

During the time of making our OPC-featured reticles, i.e., serifs, ASBs, SBs, & etc., several key process obstacles have been identified. By working jointly, we have found solution for each problem encountered. In doing so, we believe we have mobilized the entire mask making industry to work towards our goal.

For examples, one of the first problem we identified was the mask blank. The existing mask blank standard was never meant to process any OPC-featured reticles. The flatness of the quartz glass and the coating thickness for the E-beam resist were never satisfied for making our reticles. Since the need for our OPC-featured reticle started two years ago, the mask blank vendors were first caught totally un-prepared. In response to the meet the demand, they scramble to fast R/D a solution for a OPC-compatible blank. Yet it took nearly two years for the two major Japanese mask blank vendors -- Hoya & Ulcoat -- to be able to supply in a production quantity.

The domestic blank vendor, DuPont, was barely able to meet the new standard for the flatness and resist coating. But DuPont's technology for the interface cleaning between the sputtered chrome and the quartz substarte was far from ideal. As it turns out, this was the second major hurdle that the blank supplier must overcome to allow for producing decent OPC features in our reticle. Again, the two Japanese vendors found the magic and have a production formula for cemeting a very nice interface between the chrome and the quartz substrate.

As far as e-beam processing front, both our mask making vendors now have been convenienceed that using a 0.25um beam spot size is the best way for writing

our massively complicated reticles. This way most of the writing jobs can be accomplished within 3 to 5 hours write time. (As compared to spot sized 0.125um write time, it may take as many as 16 to 20 hours.) This makes e-beam writing time for our jobs become no different from the regular not-so-advanced-reticle jobs. Thus the cost can be substantially reduced. With a shorter e-beam writing time, it comes with an improved registration from one masking layer to the next.

Nevertheless, the 0.25 um beam spot writing strategy demands a very fine process control to meet our critical dimension spec. This is since the 0.25um is much coarser beam spot than the spot size of 0.125um. In our opinion, this is probably the most difficult of all problems that we experienced. It took more than a year for the Hoya Micro Mask (now called Photronics Sunnyvale) to engineer a reproducible resist and chrome etching process. Thanks to the Photronics acquisition of the HMM, now the similar process can be found both from DuPont Santa Clara and Align-Rite.

In the mean time, we have designed a full OPC featured test vehicle to help the mask vendors to dial-in their process control. The recent results shown that both vendors' process receipts are fairly matured and ready for meeting our need.

There are many other improvements made during the exercises. It will be too much to mention in detail. In any case, we have accomplished the most important goal -- that we have trained our mask vendors to be ready. And now the ball is in our court....

Fung

Hopper writes:
Fung,

The CAD folks are ready to generate masks. Here is the timeline as best we can work it out today:

Prerequisites:

1. A completely placed and routed Euterpe and Pollux (2 Euterpe and 2 Pollux on this next tapeout are planned.) We are quite close, (but not yet done), to a fully routed Euterpe that makes timing. Once we have that we can run DRC and LVS on the full chip. These may overlap. The LVS has the longer running time. It now takes about 5 days of CPU. Assume we need two runs to make sure Euterpe is clean. Assume we can overlap the last LVS with initial mask generation (if that LVS fails, we need to restart the masks).
- Geert is rebuilding Pollux now. After Pollux is rebuilt there is 4 weeks worth of work required to bring the layout up to the current level of Mobi design rules. The 4 weeks assumes 3 people in the CAD group along with help from design and mask engineers. Geert has published a detailed schedule of tasks in this area.
2. A frozen DRC flow. As a result of weekly meetings with the Fab group, we now have a substantial body of CAD implementation work ahead of us. There are several waflization, synthesis and fracture changes that need to be made. This work is not yet completed and we anticipate more changes in the coming weeks as the Fab group gathers more process data. I would estimate at least 4 weeks of CAD effort is required to make the necessary modifications we have accumulated thus far. This work is going on now.
3. About 8 weeks of CPU time to generate a full set of masks. We can generate the first mask layers almost immediately (within 1 day). The top layers (the metals) are more complex and because of several recently introduced flow changes, can no longer be run in parallel. These back-end masks will take at least 6 weeks of solid, uninterrupted CPU time. All indications are that the synthesis process will become more complex as the design rules are refined, hence the 8 week figure. [More CPU horsepower might help here, but some of this work is inherently serial.]

Therefore, if we had a routed Euterpe/Pollux today we could have the first masks out in about 8 weeks and the final mask out in about 16 weeks.

There are some caveats: We assume the availability of enough CPU cycles so that we can check, waffle, synthesize and fracture Euterpe and Pollux in parallel. We assume that our machines will not crash during this period.

We assume that the design rules do not change. Once we start this process, any change to the design rules will cause an 8 week CPU reset plus the time to verify the new rule. So that the Fab group to can further verify SDEC and iso layers, the CAD folks are going to produce a revised 150 layer of Calliope 1.

We are working on this now in parallel with other tasks.

Since the mask generation process takes substantial CPU time, we have elected to run the XOR checks (mask synthesis and fracture verification) overlapping with mask generation. If the XOR check fails, we will have to cancel a mask or masks and restart the fracture process from the beginning. The XOR checking will probably take on the same order of time as the initial synthesis and fracture, probably another 6 to 8 weeks. We should be able to run these checks in parallel (with a time-lag start) with mask generation. This will be the first time this waffle, synthesis and fracture flow has been run. We do expect some re-work. That time has not been factored in here.

-hopper

From: hopper (Mark Hofmann)
Sent: Wednesday, May 10, 1995 7:08 PM
To: fung
Cc: geert@gaea; anh@gaea; cadettes@gaea; paulp@gaea; tbr@gaea; al@charybdis; fung@charybdis; mouss@charybdis; ted@charybdis
Subject: Re: Mask vendors are ready to go...

fung writes:
Geert,

We are very happy to inform you that our mask vendors are ready to take our reticle jobs at any time now. Under our incentivized purchasing program, they are committed to ship to us with a schedule of at least one mask per day! Currently we now have two completely qualified mask vendors with a possibility of getting the third one qualified before Q4 this year. It is our goal that by the end of this year, if we choose to, we should be able to tape out three mask sets at the same time!

Since we have worked very closely with our mask vendors, they have submitted a request to us and needed an estimated answer. That is, in order to help them plan for allocating the capacity, they would like to have an estimated tape-out dates for both the Euterpe and Mnemo. It will be very thankful if you can provide the information to us as soon as you can.

To give you a very brief background: for the last three to four months, Kurt, Al, & myself have continued to work with our two major mask suppliers for perfecting the reticle making process. According to the most recent results we have seen from the two mask vendors -- DuPont (Santa Clara) & Photonics (Sunnyvale) -- we felt that the reticle processing technology have finally become much more production worthy to suit our need. The potential supplier -- Align Rite Corp. -- has come in a long way and making a very good progress. It has narrowed in their shortcomings in terms of the processing equipment. Major acquisitions will be made by the end of this month in order to have the required capability.

During the time of making our OPC-featured reticles, i.e., serifs, ASBs, SBs, & etc., several key process obstacles have been identified. By working jointly, we have found solution for each problem encountered. In doing so, we believe we have mobilized the entire mask making industry to work towards our goal.

For examples, one of the first problem we identified was the mask blank. The existing mask blank standard was never meant to process any OPC-featured reticles. The flatness of the quartz glass and the coating thickness for the E-beam resist were never satisfied for making our reticles. Since the need for our OPC-featured reticle started two years ago, the mask blank vendors were first caught totally un-prepared. In response to the meet the demand, they scramble to fast R/D a solution for a OPC-compatible blank. Yet it took nearly two years for the two major Japanese mask blank vendors -- Hoya & Ulcoat -- to be able to supply in a production quantity.

The domestic blank vendor, DuPont, was barely able to meet the new standard for the flatness and resist coating. But DuPont's technology for the interface cleaning between the sputtered chrome and the quartz substate was far from ideal. As it turns out, this was the second major hurdle that the blank supplier must overcome to allow for producing decent OPC features in our reticle. Again, the two Japanese vendors found the magic and have a production formula for cementing a very nice interface between the chrome and the quartz substrate.

As far as e-beam processing front, both our mask making vendors now have been convenience that using a 0.25um beam spot size is the best way for writing

our massively complicated reticles. This way most of the writing jobs can be accomplished within 3 to 5 hours write time. (As compared to spot sized 0.125um write time, it may take as many as 16 to 20 hours.) This makes e-beam writing time for our jobs become no different from the regular not-so-advanced-reticle jobs. Thus the cost can be substantially reduced. With a shorter e-beam writing time, it comes with a improved registration from one masking layer to the next.

Nevertheless, the 0.25 um beam spot writing strategy demands a very fine process control to meet our critical dimension spec. This is since the 0.25um is much coarser beam spot than the spot size of 0.125um. In our opinion, this is probably the most difficult of the all problems that we experienced. It took more than a year for the Hoya Micro Mask (now called Photronics Sunnyvale) to engineer a reproducible resist and chrome etching process. Thanks to the Photronics acquisition of the HMM, now the similar process can be found both from DuPont Santa Clara and Align-Rite.

In the mean time, we have designed a full OPC featured test vehicle to help the mask vendors to dial-in their process control. The recent results shown that both vendors' process recipes are fairly matured and ready for meeting our need.

They are many other improvement made during the exercises. It will be too much to mention in detail. In any case, we have accomplished the most important goal -- that we have trained our mask vendors to be ready. And now the ball is in our court....

Fung

Hopper writes:
Fung,

The CAD folks are ready to generate masks. Here is the timeline as best we can work it out today:

Prerequisites:

1. A completely placed and routed Euterpe and Pollux (2 Euterpe and 2 Pollux on this next tapeout are planned.) We are quite close, (but not yet done), to a fully routed Euterpe that makes timing. Once we have that we can run DRC and LVS on the full chip. These may overlap. The LVS has the longer running time. It now takes about 5 days of CPU. Assume we need two runs to make sure Euterpe is clean. Assume we can overlap the last LVS with initial mask generation (if that LVS fails, we need to restart the masks).
- Geert is rebuilding Pollux now. After Pollux is rebuilt there is 4 weeks worth of work required to bring the layout up to the current level of Mobi design rules. The 4 weeks assumes 3 people in the CAD group along with help from design and mask engineers. Geert has published a detailed schedule of tasks in this area.
2. A frozen DRC flow. As a result of weekly meetings with the Fab group, we now have a substantial body of CAD implementation work ahead of us. There are several wafflization, synthesis and fracture changes that need to be made. This work is not yet completed and we anticipate more changes in the coming weeks as the Fab group gathers more process data. I would estimate at least 4 weeks of CAD effort is required to make the necessary modifications we have accumulated thus far. This work is going on now.
3. About 8 weeks of CPU time to generate a full set of masks. We can generate the first mask layers almost immediately (within 1 day). The top layers (the metals) are more complex and because of several recently introduced flow changes, can no longer be run in parallel. These back-end masks will take at least 6 weeks of solid, uninterrupted CPU time. All indications are that the synthesis process will become more complex as the design rules are refined, hence the 8 week figure. [More CPU horsepower might help here, but some of this work is inherently serial.]

Therefore, if we had a routed Euterpe/Pollux today we could have the first masks out in about 8 weeks and the final mask out in about 16 weeks.

There are some caveats: We assume the availability of enough CPU cycles so that we can check, waffle, synthesize and fracture Euterpe and Pollux in parallel. We assume that our machines will not crash during this period. We assume that the design rules do not change. Once we start this process, any change to the design rules will cause an 8 week CPU reset plus the time to verify the new rule. So that the Fab group to can further verify SDEC and iso layers, the CAD folks are going to produce a revised 150 layer of Calliope 1. We are working on this now in parallel with other tasks.

Since the mask generation process takes substantial CPU time, we have elected to run the XOR checks (mask synthesis and fracture verification) overlapping with mask generation. If the XOR check fails, we will have to cancel a mask or masks and restart the fracture process from the beginning. The XOR checking will probably take on the same order of time as the initial synthesis and fracture, probably another 6 to 8 weeks. We should be able to run these checks in parallel (with a time-lag start) with mask generation. This will be the first time this waffle, synthesis and fracture flow has been run. We do expect some re-work. That time has not been factored in here.

-hopper

From: Geert Rosseel [geert@gaea.microunity.com]
Sent: Wednesday, May 10, 1995 6:37 PM
To: fung@charybdis; geert@gaea
Cc: al@charybdis; anh@gaea; cadettes@gaea; mouss@charybdis; paulp@gaea; tbr@gaea; ted@charybdis
Subject: Re: Mask vendors are ready to go...

Hi,

While we don't have a fully finished Euterpe today, we are getting very close to getting a complete Euterpe (at a slower speed than the targeted number). So, if it was decided to be worthwhile to ship a slower than 926 ps Euterpe, in principle we could do that not too long from now.

However, I am very concerned about the impact of the design rules changes that are occurring. Some of these changes require a lot of manual layout and a large amount of CAD work.

I think that freezing the design rules and tape-out flow is at this point the most critical issue in getting a Euterpe taped out.

Geert

From: geert (Geert Rosseel)
Sent: Wednesday, May 10, 1995 6:37 PM
To: fung@charybdis; geert@gaea
Cc: al@charybdis; anh@gaea; cadettes@gaea; mouss@charybdis; paulp@gaea; tbr@gaea; ted@charybdis
Subject: Re: Mask vendors are ready to go...

Hi,

While we don't have a fully finished Euterpe today, we are getting very close to getting a complete Euterpe (at a slower speed than the targeted number). So, if it was decided to be worthwhile to ship a slower than 926 ps Euterpe, in principle we could do that not too long from now.

However, I am very concerned about the impact of the design rules changes that are occurring. Some of these changes require a lot of manual layout and a large amount of CAD work.

I think that freezing the design rules and tape-out flow is at this point the most critical issue in getting a Euterpe taped out.

Geert

From: fung [fung@charybdis]
Sent: Wednesday, May 10, 1995 6:17 PM
To: Geert Rosseel
Cc: Anh Ngo; cadettes; Paul Poenisch; Tim B. Robinson; al@charybdis; fung@charybdis; mouss@charybdis; ted@charybdis
Subject: Mask vendors are ready to go...

Geert,

We are very happy to inform you that our mask vendors are ready to take our reticle jobs at any time now. Under our incentivized purchasing program, they are committed to ship to us with a schedule of at least one mask per day! Currently we now have two completely qualified mask vendors with a possibility of getting the third one qualified before Q4 this year. It is our goal that by the end of this year, if we choose to, we should be able to tape out three mask sets at the same time!

Since we have worked very closely with our mask vendors, they have submitted a request to us and needed an estimated answer. That is, in order to help them plan for allocating the capacity, they would like to have an estimated tape-out dates for both the Euterpe and Mnemo. It will be very thankful if you can provide the information to us as soon as you can.

To give you a very brief background: for the last three to four months, Kurt, Al, & myself have continued to work with our two major mask suppliers for perfecting the reticle making process. According to the most recent results we have seen from the two mask vendors -- DuPont (Santa Clara) & Photronics (Sunnyvale) -- we felt that the reticle processing technology have finally become much more production worthy to suit our need. The potential supplier -- Align Rite Corp. -- has come in a long way and making a very good progress. It has narrowed in their shortcomings in terms of the processing equipment. Major acquisitions will be made by the end of this month in order to have the required capability.

During the time of making our OPC-featured reticles, i.e., serifs, ASBs, SBs, & etc., several key process obstacles have been identified. By working jointly, we have found solution for each problem encountered. In doing so, we believe we have mobilized the entire mask making industry to work towards our goal.

For examples, one of the first problem we identified was the mask blank. The existing mask blank standard was never meant to process any OPC-featured reticles. The flatness of the quartz glass and the coating thickness for the E-beam resist were never satisfied for making our reticles. Since the need for our OPC-featured reticle started two years ago, the mask blank vendors were first caught totally un-prepared. In response to the meet the demand, they scramble to fast R/D a solution for a OPC-compatible blank. Yet it took nearly two years for the two major Japanese mask blank vendors -- Hoya & Ulcoat -- to be able to supply in a production quantity.

The domestic blank vendor, DuPont, was barely able to meet the new standard for the flatness and resist coating. But DuPont's technology for the interface cleaning between the sputtered chrome and the quartz substate was far from ideal. As it turns out, this was the second major hurdle that the blank supplier must overcome to allow for producing decent OPC features in our reticle. Again, the two Japanese vendors found the magic and have a production formula for cementing a very nice interface between the chrome and the quartz substrate.

As far as e-beam processing front, both our mask making vendors now have been convenience that using a 0.25um beam spot size is the best way for writing our massively complicated reticles. This way most of the writing jobs can be accomplished within 3 to 5 hours write time. (As compare to spot sized 0.125um write time, it may take as many as 16 to 20 hours.) This makes e-beam writing time for our jobs become no different from the regular not-so-advanced-reticle jobs. Thus the cost can be substantially reduced. With a shorter e-beam writing time, it comes with a improved registration from one masking layer to the next.

Nevertheless, the 0.25 um beam spot writing strategy demands a very fine process control to meet our critical dimension spec. This is since the 0.25um is much coarser beam spot than the spot size of 0.125um. In our opinion, this is probably the most difficult of the all problems that we experienced. It took more than a year for the Hoya Micro Mask (now called Photronics Sunnyvale) to engineer a reproducible resist and chrome etching process. Thanks to the Photronics acquisition of the HMM, now the similar process can be found both from DuPont Santa Clara and Align-Rite.

In the mean time, we have designed a full OPC featured test vehicle to help the mask vendors to dial-in their process control. The recent results shown that both vendors' process recipes are fairly matured and ready for meeting our need.

There are many other improvements made during the exercises. It will be too much to mention in detail. In any case, we have accomplished the most important goal -- that we have trained our mask vendors to be ready. And now the ball is in our court....

Fung

From: pmayer (Patricia Mayer)
Sent: Wednesday, May 10, 1995 4:55 PM
To: tbr
Cc: pmayer
Subject: PCB layout schedules

Tim,

Just wanted to let you know whats going on.

Have one more ECO for Hestia-Main today and high hopes of finishing by Monday at our 3:00 meeting. (and with a sinus infection!)

Euterpe, Mnemo and Herminator have ECO's pending with major mechanical changes. The edits are easy enough, just time consuming, like move the 80 pin connector over just a tad... Thats about 100 edits.

I foresee a design review by the middle of next week.

Howard was very receptive to doing the test boards, even offered to start the back plane and then hand it over to the next person.

I hope to move after Hestia is done, hopefully by the beginning of the week and then I need to set up the other cube for the contractor. Did we get that one? Who else should I contact for that?

-Pattie

From: ong (Warren R. Ong)
Sent: Tuesday, May 09, 1995 8:46 PM
To: vant
Cc: hopper (Mark Hofmann); geert (Geert Rosseel); tbr (Tim B. Robinson); lisar (Lisa Robinson); vo (Tom Vo); vanthof (Dave Van't Hof); wampler (Kurt Wampler); tom (Tom Laidig)
Subject: Re: euterpe lvs finished

>From vant ...

@
@
@ The latest attempt at a fullchip LVS finished and it's a tad bit cleaner, @ but not much. There still appears to be one short with the vref22 oph @ net, but I'm not too worried about that. The real problems appear to be @ inside the nb block. I've taken a quick glance at some of the locations @ and the two cells which show up immediately are:
@
@ eaffbbdh12s11x2a
@ eamemalr1wpr6x1a
@
@ There may be more. The errors are flagged inside the guts of these cells @ which could indicate a hookup problem, especially if there are clocks @ invovled. Could they be hooked up backwards?
@
@ Of the 1800 or so errors, most of them are associated with the NB block.
@ It may be necessary to run an lvs on the nb by itself (oh joy...)
@
@ The error file is:
@
@ /u/vanthof/compass/mobi/euterpe/tapeout/euterpe.compare/euterpe.lvs
@
@ Kurt is planning to look into the error file in more detail.

Glancing at the lvs report, I kind of think that the memory cells of the buffers are shorted together. The internal node of the memory cell 'collb' seems to be shorted to another node of the same name. My reasoning is (1) 'collb' of quite a few memory cells are not matched (2) later in the report, it thinks that the pld connected between vdd and this node is twice the size than the schematic size. Could it be that the memory cells are flipped (x=-x) in the array, when (I think) they should be stepped?

Back to an ealier topic:

The above two mentioned cells are the cells that I said whose layouts were probably not touched when this lvs was started.

--

Warren.

From: vanthof (vant)
Sent: Tuesday, May 09, 1995 6:32 PM
To: hopper (Mark Hofmann); geert (Geert Rosseel); tbr (Tim B. Robinson); lisar (Lisa Robinson); vo (Tom Vo)
Cc: vanthof (Dave Van't Hof); wampler (Kurt Wampler); tom (Tom Laidig); ong (Warren R. Ong)
Subject: euterpe lvs finished

The latest attempt at a fullchip LVS finished and it's a tad bit cleaner, but not much. There still appears to be one short with the vref22_0ph net, but I'm not too worried about that. The real problems appear to be inside the nb block. I've taken a quick glance at some of the locations and the two cells which show up immediately are:

```
eaffbbdh12s1lx2a  
eamemalr1wpr6x1a
```

There may be more. The errors are flagged inside the guts of these cells which could indicate a hookup problem, especially if there are clocks involved. Could they be hooked up backwards?

Of the 1800 or so errors, most of them are associated with the NB block. It may be necessary to run an lvs on the nb by itself (oh joy...).

The error file is:

```
/u/vanthof/compass/mobi/euterpe/tapeout/euterpe.compare/euterpe.lvs
```

Kurt is planning to look into the error file in more detail.

Thanks,
Dave

--
Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100
"Don't blame me! I didn't vote for him"

From: vanthof (vant)
Sent: Tuesday, May 09, 1995 6:32 PM
To: hopper (Mark Hofmann); geert (Geert Rosseel); tbr (Tim B. Robinson); lisar (Lisa Robinson); vo (Tom Vo)
Cc: vanthof (Dave Van't Hof); wampler (Kurt Wampler); tom (Tom Laidig); ong (Warren R. Ong)
Subject: euterpe lvs finished

The latest attempt at a fullchip LVS finished and it's a tad bit cleaner, but not much. There still appears to be one short with the vref22_0ph net, but I'm not too worried about that. The real problems appear to be inside the nb block. I've taken a quick glance at some of the locations and the two cells which show up immediately are:

```
eaffbbdh12s1lx2a  
eamemalrlwpr6x1a
```

There may be more. The errors are flagged inside the guts of these cells which could indicate a hookup problem, especially if there are clocks involved. Could they be hooked up backwards?

Of the 1800 or so errors, most of them are associated with the NB block. It may be necessary to run an lvs on the nb by itself (oh joy...).

The error file is:

```
/u/vanthof/compass/mobi/euterpe/tapeout/euterpe.compare/euterpe.lvs
```

Kurt is planning to look into the error file in more detail.

Thanks,
Dave

--

Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100
"Don't blame me! I didn't vote for him"

From: geert (Geert Rosseel)
Sent: Tuesday, May 09, 1995 2:05 PM
To: bill; bpw; briant; dtacmo; efelias; two; geert; hopper; lisar; mikew; mudge; ong; orlando; stick; tau; tbr; vanthof; vikki; wampler; warren; wingard
Subject: Summary of Cronus/Atlas Meeting

Hi,

Here is a summary of last weeks Cronus/Atlas Status meeting

1. Baseplate

Current status : We have a gards-compiled cronus baseplate which can be used for routing experiments,

Plan : By May 31, we want to have a final baseplate.

Action items :

- * Decide on final die size : Drew
- * Pading assignment : Warren
- * Floorplan, custom block placements : Warren
- * Clock Tree generation : Kurt, Bill
- * Build dummy cells for
TTL I/O and : B.P
IOBYTE : Dane

2. Custom blocks

Current status : GTLB done
Register file : end of this week
Caches : layout at top-level
Tags : layout has not started yet
NB : layout has just started
TTL I/O : not designed yet
IOBYTE : not started yet

Plan : finish all custom blocks by end of May
except TTL I/O and IOBYTE

Action Items :

- * Register File : B.P, Mike, Orlando
- * Caches : Bruce, Efelias, Dtacmo
- * Tags : Bruce, Efelias, Dtacmo
- * NB : Vikki, B.P.
- * TTL I/O : B.P., ??
- * IOBYTE layout can start by end of next week.
We'll need schematics by then : Dane, Bill

3. Atlas database

Current status : Database builds from toplevel. An initial version of about everything is there.

Plan : Have a complete and accurate Atlas database by the end of May
In building Cronus sub-blocks, everybody should point to /u/chip/atlas and not local versions. That will catch database problems early.

Action Items :

- * Review XL, XS cells

- timing, layout, cap models : Warren
- * Finish timing numbers : Fred
- * Complete verilog libraries : Brianl
- * Complete Synopsis libraries : Brianl

4. Cronus/verilog/bsrc, Makefile.rules, GARDS

Current status : Initial versions of Makefile.rules exist. We can build a sub-block using the atlas database on the checked in baseplate model.

Plan : Have a "good" Makefile.rules that works well enough that other people can start mapping Euterpe blocks by the end of May

Implement " put the top-level together" strategy in the equivalent of euterpe/verilog/bsrc/Makefile.tst

Action Items :

- * Makefile.rules : Brianl
- * GARDS model additions to deal with power and clock obs : Tau, Kurt
- * Pim2pif upgrade : Hopper
- * Take euterpe snapshot : Tbr
- * Floorplanning tools : Drew
- * verilog/bsrc/Makefile : Drew

5. "Implementation of mapping"

No plans for this month

6. Packaging

Current status : We have "a plan" (read all about it in Mosaic)

Plan : By the end of May, make decisions on all aspects of "the plan"

Action Items :

- * Call a set of meetings during May to turn our plan into a set of decisions : Geert

7. Test

Current Status : Mudge & Mark Warren have taken the responsibility for Cronus test, both die sort and test of packaged parts,

Plan : Same as packaging : by the end of May, we want to decide on our plan so that only implementation issues are left.

Action Items : Johnny and Mark own this, so they can decide.

From: hopper (Mark Hofmann)
Sent: Tuesday, May 09, 1995 1:16 PM
To: Kurt Wampler
Cc: geert (Geert Rosseel); tbr (Tim B. Robinson); wampler (Kurt Wampler)
Subject: Re: GARDS 7.117?

Kurt Wampler writes:

Last Friday we received a fixed GPLACE which closes out the two TARs we had registered against it. I have tested it with the original cases we submitted on the bug tape and it worked correctly. SVR will be pressing for us to accept the 7.117 release soon so they can make the source escrow of it for us.

When would be a good time to switch over to the 7.117 release for general use? There may still be other bugs lurking in this release that we haven't yet uncovered, but it has received a moderate amount of exercise by Brian, Drew, and me.

The PGRROUTE fix is still a week or two away. It would be good to defer the capture of the escrow tape until it includes the fixed PGRROUTE.

Comments?

- Kurt

I agree that it would be good to wait for the PGRROUTE fix. It's a big deal to us and they appear to be close to a final release.

-hopper

From: wampler (Kurt Wampler)
Sent: Tuesday, May 09, 1995 12:53 PM
To: geert; tbr
Cc: hopper; wampler
Subject: GARDS 7.117?

Last Friday we received a fixed GPLACE which closes out the two TARs we had registered against it. I have tested it with the original cases we submitted on the bug tape and it worked correctly. SVR will be pressing for us to accept the 7.117 release soon so they can make the source escrow of it for us.

When would be a good time to switch over to the 7.117 release for general use? There may still be other bugs lurking in this release that we haven't yet uncovered, but it has received a moderate amount of exercise by Brian, Drew, and me.

The PGROUTE fix is still a week or two away. It would be good to defer the capture of the escrow tape until it includes the fixed PGROUTE.

Comments?

- Kurt

From: vanthof (vant)
Sent: Tuesday, May 09, 1995 1:30 AM
To: tom (Tom Laidig); wampler (Kurt Wampler); paulp (Paul Poenisch)
Cc: vanthof (Dave Van't Hof); tbr (Tim B. Robinson); geert (Geert Rosseel); hopper (Mark Hofmann)
Subject: changes to the drc flow

I've been trying to implement the latest preliminary revision of the drc specs and I'm having some major problems. The inter-relationships between all metals and ABS's are so complex and intricate verifying they are correct may not be possible. The problems I'm having are with the footnotes for the XX.c rules (min spacing except when metals are larger than 2.5 microns, then .25 microns will be cut away automatically, except around 'almost' min connecting layers). In order to verify the rule, I must code in a potential synthesis algorithm for the metals/vias. Unfortunately this can never be 100% accurate as I don't have any way to synthesize the ABSs as they would be at tapeout. In addition, this set of rules has probably doubled the run time and significantly increased the memory usage for drc runs.

When the rules used to state explicitly that metals over a certain width must have a certain spacing, that was a very easy (in comparison) rule to check across all layers.

I have a drc flow which is a first attempt at checking the XX.c class of rules across all metals, however, I'm fairly certain it's not 100% complete for the air-bridged metals. These layers are the nasty ones.

For the record, I currently haven't put any thought into checking the XX.d rules (min sized hole is 30 udrs, it will be a toughy) and the XX.e and XX.f rules (min and max area coverage) have not been implemented yet. These three new rules are also going to increase the runtimes a bunch.

I'll keep working on this, and it may get better after I think on it some more, but every time I do look at this, it gets more complex and more difficult to verify it's done correct.

Dave

--

Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@micrunity.com 1 408 734-8100
"Don't blame me! I didn't vote for him"

From: vanthof (vant)
Sent: Tuesday, May 09, 1995 1:30 AM
To: tom (Tom Laidig); wampler (Kurt Wampler); paulp (Paul Poenisch)
Cc: vanthof (Dave Van't Hof); tbr (Tim B. Robinson); geert (Geert Rosseel); hopper (Mark Hofmann)
Subject: changes to the drc flow

I've been trying to implement the latest preliminary revision of the drc specs and I'm having some major problems. The inter-relationships between all metals and ABS's are so complex and intricate verifying they are correct may not be possible. The problems I'm having are with the footnotes for the XX.c rules (min spacing except when metals are larger than 2.5 microns, then .25 microns will be cut away automatically, except around 'almost' min connecting layers). In order to verify the rule, I must code in a potential synthesis algorithm for the metals/vias. Unfortunately this can never be 100% accurate as I don't have any way to synthesize the ABSs as they would be at tapeout. In addition, this set of rules has probably doubled the run time and significantly increased the memory usage for drc runs.

When the rules used to state explicitly that metals over a certain width must have a certain spacing, that was a very easy (in comparison) rule to check across all layers.

I have a drc flow which is a first attempt at checking the XX.c class of rules across all metals, however, I'm fairly certain it's not 100% complete for the air-bridged metals. These layers are the nasty ones.

For the record, I currently haven't put any thought into checking the XX.d rules (min sized hole is 30 udrs, it will be a toughy) and the XX.e and XX.f rules (min and max area coverage) have not been implemented yet. These three new rules are also going to increase the runtimes a bunch.

I'll keep working on this, and it may get better after I think on it some more, but every time I do look at this, it gets more complex and more difficult to verify it's done correct.

Dave

--

Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100
"Don't blame me! I didn't vote for him"

From: lisar (Lisa Robinson)
Sent: Monday, May 08, 1995 12:20 PM
To: tbr; warren; wayne
Cc: geert; mudge; jeffm; pandora; euterpe
Subject: euterpe/cronus test boards

Hardware debug stations

I'm not sure if I have addressed the email to all of those that are interested so please if you or someone you know is interested please join the meeting.

Some background.

The build plan for Pandora Systems includes stations for Chip Characterizations. These stations would likely live in the lab close to the test group and other chip characterization equipment.

My thoughts were that the Pandora Modules (euterpe/mnemo) are well suited to carry packaged parts for debug since few additional components are required and the cost per board would be fairly low (by comparison to a custom socket or fixture). In addition the parts could be run at speed.

There is also a requirement to mount unbumped Cronus parts in a test fixture (not run at speed).

I'd like to have a meeting to discuss the "test board" strategy.

Wednesday at 3.00pm.

Lisa R.

From: William Herndon [bill@polyhymnia.microunity.com]
Sent: Thursday, May 04, 1995 12:02 PM
To: hopper@polyhymnia.microunity.com
Cc: Geert Rosseel; geert@polyhymnia.microunity.com; graham@polyhymnia.microunity.com; lisar@polyhymnia.microunity.com; mudge@polyhymnia.microunity.com; paulp@polyhymnia.microunity.com; rich@polyhymnia.microunity.com; tau@polyhymnia.microunity.com; tbr@polyhymnia.microunity.com; tomho@polyhymnia.microunity.com; vanthof@polyhymnia.microunity.com
Subject: Re: Status of chips : Summary

There was discussion of some disciplines in design/definition of castor must be followed. Geert mentions a requestor and test plan requirement. tom mentioned no intermingling of test sites. What both are saying is that it has to be reasonably straightforward and well thought out. It is clear that we need basic device information that is on castor. We also need basic information on metal continuity and lack of shorts. It seems it ought to be easy to make a test that says m1 is continuous, m2 is continuous, m1 via m2 path is continuous, m1 does not short to m2 when there aren't vias, and then repeat m2 m3, m3 m4.

On Thu, 4 May 1995 hopper@polyhymnia.microunity.com wrote:

> Geert Rosseel writes:

> [snip]

> *****
> PROPOSED PLAN OF ACTION
> *****

> 5. On the following tape-out = Mnemo , we will have two sites
> of the reticle taken up by Mnemo and two sites by Castor
> The Mnemo tape-out date will determine the tape-out date
> of this reticle. If we have not completely finished
> Castor, so be it.

> I do believe that a good version of Castor is very important to the fab.
> Every effort should be made to have a finished Castor to tapeout with Mnemo.
> One of the reasons that Castor requires so much re-work now is that it
> was not thoroughly checked over in its first incarnation. If we are
> going to commit the resources to re-doing Castor, then I think we
> should try hard not to make that mistake this time around.

> Both Pollux and Castor are modular designs, so we can
> tape-out partial chips.

> -hopper

From: hopper (Mark Hofmann)
Sent: Thursday, May 04, 1995 8:35 AM
To: Geert Rosseel
Cc: bill (William Herndon); geert (Geert Rosseel); graham (Graham Y. Mostyn); lisar (Lisa Robinson); mudge (John Mudge); paulp (Paul Poenisch); rich (Rich McCauley); tau; tbr (Tim B. Robinson); tomho (Tom Ho); vanthof (Dave Van't Hof)
Subject: Re: Status of chips : Summary

Geert Rosseel writes:

[snip]

PROPOSED PLAN OF ACTION

5. On the following tape-out = Mnemo , we will have two sites of the reticle taken up by Mnemo and two sites by Castor. The Mnemo tape-out date will determine the tape-out date of this reticle. If we have not completely finished Castor, so be it.

I do believe that a good version of Castor is very important to the fab. Every effort should be made to have a finished Castor to tapeout with Mnemo. One of the reasons that Castor requires so much re-work now is that it was not thoroughly checked over in its first incarnation. If we are going to commit the resources to re-doing Castor, then I think we should try hard not to make that mistake this time around.

Both Pollux and Castor are modular designs, so we can tape-out partial chips.

-hopper

From: geert (Geert Rosseel)
Sent: Wednesday, May 03, 1995 11:57 PM
To: bill; geert; graham; hopper; lisar; mudge; paulp; rich; tau; tbr; tomho; vanthof
Subject: Status of chips : Summary

Hi,

Here is a summary of this afternoon's meeting on the designs currently in the fab :

CURRENT STATUS

	major problems	expected yield	Effort required to fix (# months)
CASTOR :	metals REALLY bad	0	2-3 months, 3 people
POLLUX :	metals bad	0	1-2 months, 2 people
CALLIOPE0:	metals, padding bad	0	2 months, 2 people
ORCHIS:	padding bad	0 to very low	1 month, 2 people
CALLIOPE1	padding bad	0 to very low	1 month, 2 people

Here is the input from everybody in the room :

CASTOR :
Johnny's group needs CASTOR for parametric test.
Without Castor, no parametrics. The parametrics are essential for debugging all analog circuits.

POLLUX :
All circuit-designers believe POLLUX will be essential to debug our circuits, especially the analog circuits.

CALLIOPE0 :
Its initial purpose was to help us develop the CAD flow to tape-out a "real" design.
Calliope0 has served this purpose very well, but at this point, nobody has any strong need for this chip.

ORCHIS :
Useful as a yield vehicle once the process is up.
The pad-ring, seal-ring structures are not very good, and need work (this is mostly the work we have to do to fix the euterpe/mnemo pads). As is, the yield is expected to be very low and the lifetime of the part after air-bridge could be as low as 30 minutes.

CALLIOPE1 :
Has the same problems as Orchis.

PROPOSED PLAN OF ACTION

1. No work will be done on Calliope0.
2. We will fix up Pollux, this is a FULL MASK SET REVISION

3. We will fix up Castor, also a FULL MASK SET REVISION
Castor will be redone so that it not only conforms
to the current design rules, but also is compatible with our design
methodology.
Geert (editors comment) feels strongly that any test-circuit that
we implement should :

- > Have a "requester" , if we have something on the current
Castor that nobody claims ownership for, it goes ...
- > Be backed up by a test-plan (if there is no plan to test it
by now , 16 months after initial tape-out, it's probably
not an important circuit)

4. On the first next tape-out = Euterpe, we will have two sites
of the reticle taken up by Euterpe and two sites by Pollux
(We want two of each for die to die inspection)

5. On the following tape-out = Mnemo , we will have two sites
of the reticle taken up by Mnemo and two sites by Castor
The Mnemo tape-out date will determine the tape-out date
of this reticle. If we have not completely finished
Castor, so be it.

Both Pollux and Castor are modular designs, so we can
tape-out partial chips.

ACTION ITEMS

We will have separate small meetings on Castor and Pollux to find out in greater detail
what exactly needs to be done and how much effort it will take.

Once we have that data, we should review our current plan (change it if necessary) and
decide on how to implement these changes with minimal impact on the Euterpe, Mnemo and
Cronus schedules.

Geert

From: doi (Derek Iverson)
Sent: Wednesday, May 03, 1995 5:25 PM
To: guarino; gmo; gregg; claseman; lisa; jeffm
Cc: hestla
Subject: Software Bringup Meeting Minutes - May 3, 1995

Software Bringup Meeting

May 3, 1995

Next Meeting: May 10 at 2:00 pm.

Attendees: guarino, gmo, doi, gregg, claseman, lisa, jeffm

New Action Items

Item: When do we have a full calliope simulation available (IKOS)?

Who: lisar

Status: New

04/26 This topic was raised as we talked about when the Snap/Unsnap item should be brought back from the Suspended list.

Review of Action Items

Item: Modify startup code in stb/stand to read the cerberus node number

Who: gmo

Status: Pending

05/03 No new progress

Item: Plan for testing remote debugging environment

Who: everyone

Status: Pending

04/13 This will be discussed at the next meeting.

04/20 Short discussion on the pieces required and if we want to have a standalone version of the hostio software. Wally has started on this already.

04/26 Any work with snoopy/dram models will begin after euterpe tapeout when jeff has more time.

05/03 Wayne F. expects to have a prototype ISA card ready for testing tomorrow. Gmo is in the process of writing code to test the card.

Item: Schedule a meeting to go over how to run stuff on HW simulator and read HW traces.

Who: jeffm, lisar, doi

Status: Pending

04/05 doi is going to talk to lisar about when this would be appropriate. Lisar has volunteered to show up at the next meeting and discuss this.

04/13 lisar presented a strategy that will give other the ability to run and do software debug with the hardware simulators. Jeffm to write up a crib sheet that will aid with debugging in

the hardware accelerator environment.

04/20 Dave Tomlinson will require such a session. Sometime after May 8.

04/26 Suspended until 05/08.

Item: ukernel needs to detect machine checks and 'do the right thing'

Who: lisa

Status: Pending

04/13 No new progress.

04/20 On list of things to do. Lower priority.

04/26 Ditto.

05/03 Ditto.

Item: Modify tests in diag tree to use tcc instead of tgcc

Who: guarino, jeffm, doi

Status: Pending

03/29 Loretta has changed the diag tests to use tcc instead

of tgcc but has not checked in the changes.

Lisar wants these changes to be made after we are able to run the tests successfully using tgcc.

04/20 Loretta has the stand/diag tests converted to use TCC.

Derek is going to modify the Makefiles in stand/diag/memhi so that the programs that required tgcc will have an explicit rule.

04/26 Loretta is ready to check in the TCC related changes. Derek is

to check with lisar if this is OK.

05/03 Loretta was ready to check in the changes but a last minute test uncovered a compiler bug. Checkin will happen after the tests build and run OK.

Suspended Items

Item: Unsnap code

Who: sandeep, guarino

Status: Suspended.

02/15 The issue of restarting the hardware from an IKOS dump was discussed and the need for an architectural snap/unsnap facility was questioned.

Since the meeting it has been re-discovered (jeffm wasn't there to remind us of an earlier decision) that we are planning on loading architectural state into an IKOS simulation and not from a total IKOS logic dump.

We also determined that when it came time to run some of the larger tests (real-time benchmark) we would need the capability to start an IKOS simulation from an architectural dump anyhow.

03/01 For the short term we are going to focus on a simpler approach for loading and running DVTs, the kernel, and kernel tests. This item will likely come back in April.

Item: Create performance test plan

Who: jeffm, guarino

Status: [11/30] No progress as focus is on functionality.

We continue to run tests to help us compare terp vs hardware performance.

We still need to put together the actual performance tests that need to be run on the hardware.

Completed Items

None.

Terp Feature Status

inprog o Add support for host I/O through the sdram
inprog o Holes in address spaces => machine checks, exceptions, etc.
done o Reflect "forward progress" change in the hardware
 - believed done.
inprog o Ifetch protection granularity
 - performance vrs accuracy tradeoff
inprog o Fetch instructions as octlets
done o Simulate Ifetch queue
 o Accuray wrt HW simulator(s?)
 o Better latency model for Calliope accesses
 o Implement hardware configuration through Cerberus regs
 (SDRAM paramters, dram enable?)
 o Checkpoints/Snapshots
 o hermes and cerberus timeout machine checks
 - does jeff have a test for this? Currently the TSA is
 followed (immediate mchk) instead of waiting for watchdog
 timeout.
inprog o ability for terp to load hermes sections
 - lisar would like this functionality added

Performance Test Status

The hermes, rom, and dcache tests have been checked in and run on the Zycad simulator.
Tim is going to publish these numbers in a posting to 'euterpe'.
Lisa is going to sent Tim a list of the timing numbers that would be most useful for terp.
Tim is working with khp to do some timing of critical loops.

Test Status and General Discussion

The nullTest *still* hasn't run successfully on the IKOS. The latest run has found a bug in the HW (the only 'bug' that is being worked at the moment). We are in the process of getting a trace to better understand the problem.

There is a growing list of tests that need to be written and run before tapeout. This list currently consists of about 25 tests.

From: tbr
Sent: Wednesday, May 03, 1995 4:47 PM
To: wampler (Kurt Wampler)
Subject: Re: Horror story

Kurt Wampler wrote (on Wed May 3):

>I'm well aware of the issue! However, to the extent that anything
>mission critical is on the auspex, our existing system of rotating the
>last but one full backup through the bank vault should be fairly
>effective. A more comprehensive disaster recovery plan would be on his
>agenda since obviously there are a lot more issues than just that.

>I have to admit though that since Eric left I have so far been unable
>to get over all the administrative hurdles necessary for anyone else
>to get access to the vault, though that should be closed out this week.

I assume you were aware that I have both a key to the vault room and the vault combination. The key was issued by Tom Buckmaster (for access to the reticle tape cabinet) and the combination was given to me by Eric M. (The SVR escrow package is kept in the vault.)

Yes. By "vault" I was referring to the box we have at the bank vault at BofA in sunnyvale where the offsite copies are held.

>As regards a more general facility (say to archive copies of reticle
>tapes) we would have to address the fact that mouss does not want that
>in the hands of any third party.

Getting a copy of the reticle tapes offsite in a secure place is equally important. None of them are even in the fire safe at this point due to lack of space. Also, the "design snapshot" tapes right now are kept in a drawer in Tom L.'s desk. If a fire were to hit, we could lose our entire ability to reconstruct the reticles.

Those should get moved to the safe at least. However, (and maybe this has changed because of temporary lack of auspex space) don't we have the snapshots all still on line? If so copies should be in our full backups, and therefore hed offsite. I would assume from them we could regenerate the reticle tapes.

Let me know if there's any way I can help substantiate the need for offsite storage.

Thanks. Another issue that concerns me a little is the need to read old tapes periodically to keep them in good condition. It's no use having the data offsite if we cannot read it when we need to.

Tim

From: tbr
Sent: Wednesday, May 03, 1995 3:45 PM
To: wampler (Kurt Wampler)
Subject: Horror story

Kurt Wampler wrote (on Tue May 2):

Just came across this posting; a cautionary tale of someone who had no off-site tape storage and was relying on their fire safe to preserve their backup tapes. I hope we can add the establishment of a real off-site tape storage facility to Bob Van Cleef's "to do" list.

I'm well aware of the issue! However, to the extent that anything mission critical is on the auspex, our existing system of rotating the last but one full backup through the bank vault should be fairly effective. A more comprehensive disaster recovery plan would be on his agenda since obviously there are a lot more issues than just that.

I have to admit though that since Eric left I have so far been unable to get over all the administrative hurdles necessary for anyone else to get access to the vault, though that should be closed out this week.

As regards a more general facility (say to archive copies of reticle tapes) we would have to address the fact that Mouss does not want that in the hands of any third party.

Tim

From: paulp (Paul Poenisch)
Sent: Wednesday, May 03, 1995 10:54 AM
To: Tom Laidig [tau]
Cc: vanthof (Dave Van't Hof); wampler (Kurt Wampler); geert (Geert Rosseel); al (Albert Matthews); paulp (Paul Poenisch)
Subject: Re: new metal rule

Tom Laidig wrote:

>
> Paul Poenisch writes:

>
> After reviewing the pictures I have of the structures that are giving
> us problems in the fab during lift-off I've concluded that there
> needs to be a rule about coincident edges. I would like to propose
> the following rule which would be included in the layers from Contact
> Pedestal through Metal 4 excluding the ABS layers (it's not needed on
> Via 45 because of the way it's written).

>
> For layer X:

> Minimum space between features when edges on both sides
> are coincident within 0.25 um of edges on on layer X+1 0.75 um

> You mean "... on both sides [of the space] are coincident...", right?

Yes.

>
>
> As written this rule will possibly cause two problems:
> 1. Shadowing of Metal 4 by Via 45 would cause many violations, it may be
> allowable to waive this rule for metal 4 but I'm not sure.
> 2. When there are two adjacent stacks of via/metal features reaching from
> an upper layer to a lower layer (ie. metal 4 to metal 2) this rule might
> be violated (the expanding of vias in situations like this might get by
> though).

>
> I'm not sure what other problems there might be with this rule or how
> hard it would be to check. Please think about it and let me know
> what your thoughts are.

> I've been putting off responding to this, but here goes...

> This rule frankly scares the willies out of me. Dave thinks he has
> some ideas for how to check it, but none of us seems to have any ideas
> on how to obey it.

>
> Certainly none of the existing hand layout (pads, caches, register
> file, you name it) comes even close to meeting this rule. I don't
> know how much effort it would be to change all this, but the time
> units will be months.

>
> For auto-routing, we can simply tell the router that vias can't be
> placed next to each other. The routing won't fit, but we won't
> violate the rule.

>
> In cmos sofa cells, we can not use the track next to vdd or vss
> to contact a transistor -- this is currently done often.


>
> I spot-checked the ecl leaf cells, and I'm pretty confident that I
> can coax leafmold to build about half of them successfully with this
> rule (not counting contact pedestal; see next item).

> It is not possible to simultaneously contact the emitter and base
> of any bjt in the ecl atom, except by contacting a corner of the
> base (where the min spacing is corner-corner).

> It is not possible to contact the source and drain of a
> minimum-length mosfet without snaking out in SDEC or sliming contact
> pedestal away from the gate to where it can be contacted in metal.

> A quick eyeball look suggests to me that we might only experience a
> 12% increase in the size of the cache, but I didn't take into
> account the poly wafflization rules, which will probably force that
> increment up.

> This is just a quick list; I'm sure it's only the tip of the iceberg.
> My wild estimate is that in 6-9 months we could have a 12x12 mm
> euterpe in the state we thought we had euterpe in this morning.

> --
> 
>
OK. It seems as if the rule in it's present form is too limiting. I suspected that this might be the case. There are two possible modifications to the rule which I think would relieve a lot of the problems, they take significantly different approaches.

The simplest change would be to allow two layer edge stackups but prohibit three layer edge stacks. I think this would get rid of the transistor related problems but might still require some changes in the power buss scheme. The problem with this approach is that I don't know if it will work. To date I've seen the problem this rule is ment to address when three layers have stacked up but I haven't caught a lot with just two layers in place yet. I'll try to do that this morning but it will depend on exactly where lots are in the line.

A somewhat harder to check version of the rule would allow stacking of edges for some short distance but prohibit it for longer stretches, say 5 um. The problem with this rule, besides the more complex DRC program, is that a small notch in one of the layers (0.25 um by 0.5 um) would meet the rule but probably not fix the problem. This means that the conditions of the rule would have to be tightened in other ways (such as increasing the window of positions that are considered stacked edges) again making the rule troublesome to the current layout.

I would like to discuss this some this morning at 10:00.

Paul.

From: vanthof (vant)
Sent: Wednesday, May 03, 1995 8:22 AM
To: orlando (Orlando Hernando); geert (Geert Rosseel)
Cc: vanthof (Dave Van't Hof); hopper (Mark Hofmann)
Subject: euterpep drc's

The euterpep drc run is only about 7kb in size, but there are some errors to look at. In particular, the metal S1 and metal S2 max spacing errors.

The error file is:

/u/vanthof/compass/mobi/euterpe/tapeout/euterpep.err

Dave

--

Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100
"Don't blame me! I didn't vote for him"

From: tom (Tom Laidig [tau])
Sent: Tuesday, May 02, 1995 7:33 PM
To: Paul Poenisch
Cc: vanthof (Dave Van't Hof); wampler (Kurt Wampler); al (Albert Matthews); tau; geert (Geert Rosseel)
Subject: Re: new metal rule

Paul Poenisch writes:

After reviewing the pictures I have of the structures that are giving us problems in the fab during lift-off I've concluded that there needs to be a rule about coincident edges. I would like to propose the following rule which would be included in the layers from Contact Pedestal through Metal 4 excluding the ABS layers (it's not needed on Via 45 because of the way it's written).

For layer X:

Minimum space between features when edges on both sides
are coincident within 0.25 um of edges on on layer X+1 0.75 um

You mean "... on both sides [of the space] are coincident...", right?

As written this rule will possibly cause two problems:

1. Shadowing of Metal 4 by Via 45 would cause many violations, it may be allowable to waive this rule for metal 4 but I'm not sure.
2. When there are two adjacent stacks of via/metal features reaching from an upper layer to a lower layer (ie. metal 4 to metal 2) this rule might be violated (the expanding of vias in situations like this might get by though).

I'm not sure what other problems there might be with this rule or how hard it would be to check. Please think about it and let me know what your thoughts are.

I've been putting off responding to this, but here goes...

This rule frankly scares the willies out of me. Dave thinks he has some ideas for how to check it, but none of us seems to have any ideas on how to obey it.

Certainly none of the existing hand layout (pads, caches, register file, you name it) comes even close to meeting this rule. I don't know how much effort it would be to change all this, but the time units will be months.

For auto-routing, we can simply tell the router that vias can't be placed next to each other. The routing won't fit, but we won't violate the rule.

In cmos sofa cells, we can not use the track next to vdd or vss to contact a transistor -- this is currently done often.

I spot-checked the ecl leaf cells, and I'm pretty confident that I can coax leafmold to build about half of them successfully with this rule (not counting contact pedestal; see next item).

It is not possible to simultaneously contact the emitter and base of any bjt in the ecl atom, except by contacting a corner of the base (where the min spacing is corner-corner).

It is not possible to contact the source and drain of a

minimum-length mosfet without snaking out in SDEC or slimming contact pedestal away from the gate to where it can be contacted in metal.

A quick eyeball look suggests to me that we might only experience a 12% increase in the size of the cache, but I didn't take into account the poly wafflization rules, which will probably force that increment up.

This is just a quick list; I'm sure it's only the tip of the iceberg. My wild estimate is that in 6-9 months we could have a 12x12 mm euterpe in the state we thought we had euterpe in this morning.

--

'\`

From: tbr
Sent: Tuesday, May 02, 1995 6:08 PM
To: graham (Graham Y. Mostyn)
Cc: graham
Subject: Need to strategize mask changes

Graham Y. Mostyn wrote (on Fri Apr 28):

Tim, I believe we need to hold a strategy meeting to discuss the selection and timing of mask revisions for Calliope0, Calliope1 and Castor/Pollux. (An important part of the equation is that the CAD portion of these activities should be coordinated with - and not adversely impact - CAD activities for the Euterpe/Mnemo tapeout. While circuit design work could start now, I see a bottleneck at the CAD stage later.)

Before calling such a meeting, however, I'd like to hear your thoughts on the following. Perhaps this session should be followed with a broader-scope discussion on scheduling and prioritization of all MUSE mask set tapeouts.

Objective: Define and schedule Pollux and Castor mask changes.
Invitees: paulp, mudge, hopper, bill, geert, + staff

Since scheduling would be so important, we need to have lisar involved.

- Select DRCs for correction on Pollux (metal and perhaps diffusion).
- Provision of SOFA ring oscillator on Pollux, requested by fab.

We already have SOFA ring oscillators, both CMOS and ECL.

- Addition of new metal waffle rules to Pollux and Castor to allow manufacture.
- Consider necessary changes to Castor

>From informal discussions I've had with paul, it's clear that Castor is a major problem.

- Schedule of circuit design work, CAD activities and tapeout; relation to Euterpe/Mnemo mask scheduling.

Is there really circuit design work necessary? My understanding was that what we needed were layout changes driven by the changes in the design rules.

I note the above list only covers castor/pollux and not the Calliope variants. However, I think it is extremely important we consider very carefully the minimum set of things we need to do to get the essential material we need to debug the process. We must not allow the door to be opened to whole sale redesign. For that reason I favor restricting the scope to castor and pollux with perhaps a limited Calliope0 discussion restricted to how to make the current version safe to the process if we choose to make changes only to the metal layers.

Tim

From: hopper (Mark Hofmann)
Sent: Tuesday, May 02, 1995 5:34 PM
To: vant
Cc: geert (Geert Rosseel); lisar (Lisa Robinson); vo (Tom Vo); tbr (Tim B. Robinson); vanthof (Dave Van't Hof); wampler (Kurt Wampler); tom (Tom Laidig)
Subject: Re: euterpe lvs finished

vant writes:

Well, the latest euterpe lvs finished and here's the results:

NUMBER OF UN-MATCHED SCHEMATICS DEVICES	=	925
NUMBER OF UN-MATCHED LAYOUT DEVICES	=	788
NUMBER OF MATCHED SCHEMATICS DEVICES	=	2095418
NUMBER OF MATCHED LAYOUT DEVICES	=	2095418

Not too bad. There are still a lot of opens, but this time they don't appear to be related to the vref lines, but something to do with columns???

Anyway, the output file is:

/u/vanthof/compass/mobi/euterpe/tapeout/euterpe.compare/euterpe.lvs

And it's only 5.5MB this time! mucho better.

I'll keep looking at it, but if anyone else is interested, please do have a gander.

Thanks, Dave.

Sounds like the output is at least a manageable size this time.

-hopper

From: vanthof (vant)
Sent: Tuesday, May 02, 1995 5:32 PM
To: geert (Geert Rosseel); hopper (Mark Hofmann); lisar (Lisa Robinson); vo (Tom Vo); tbr (Tim B. Robinson)
Cc: vanthof (Dave Van't Hof); wampler (Kurt Wampler); tom (Tom Laidig)
Subject: euterpe lvs finished

Well, the latest euterpe lvs finished and here's the results:

NUMBER OF UN-MATCHED SCHEMATICS DEVICES	=	925
NUMBER OF UN-MATCHED LAYOUT DEVICES	=	788
NUMBER OF MATCHED SCHEMATICS DEVICES	=	2095418
NUMBER OF MATCHED LAYOUT DEVICES	=	2095418

Not too bad. There are still a lot of opens, but this time they don't appear to be related to the vref lines, but something to do with columns???

Anyway, the output file is:

`/u/vanthof/compass/mobi/euterpe/tapeout/euterpe.compare/euterpe.lvs`

And it's only 5.5MB this time! mucho better.

I'll keep looking at it, but if anyone else is interested, please do have a gander.

Dave

--

Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100
"Don't blame me! I didn't vote for him"

From: bobm (Bob Morgan)
Sent: Tuesday, May 02, 1995 11:46 AM
To: tbr; abbott; gmo; craig; mws; dickson; woody; billz; jeffm
Subject: Euterpe review draft

Hi everybody,
I've distributed review copies of the new and improved (hopefully) Euterpe MicroArchitecture book. Curtis and Craig, I'll deliver yours when you get in, or leave it in your mail slot.

I know there's not enough time to give it a thorough review today, but it's been through that once before, so I hope that won't be as much of an issue.

There are still a few technical items on my list of things to do. But there's been some pretty big reorganization, and I wanted to let you all see it before too much time went by.

I don't think trying to get ten people in a room to discuss this is a very wise move, so this is how I'd like to do it. I'd like to get the style reviewers, Tim, Curtis, Guillermo, and Craig, to look over the document as much as they can before 2:00. Then I'd like to meet at 2:00 to talk briefly about what you think of the organization of the book.

For the technical reviewers, I'd like you to take a look at the organization of the book too. If you have a chance to do it before 2:00, I'd appreciate you sending me a note, or drop by my office, with any comments you have. As I said, there are still a few technical changes I need to make, and I'd like to continue receiving any comments or corrections you have over the next couple weeks before we send this out to "outsiders" (oooh).

Thanks to you all for your help. Sorry about the length of this note.
Tech writers are supposed to value brevity.
Bob

From: paulp@microunity.com
Sent: Monday, May 01, 1995 5:40 PM
To: paulp@rhea
Cc: geert@rhea; hopper@rhea; tbr@rhea; fwo@rhea
Subject: metalcells/2265: Crack protection ring waffle pattern

>Number: 2265
>Category: metalcells
>Synopsis: Crack protection ring waffle pattern
>Confidential: no
>Severity: critical
>Priority: high
>Responsible: paulp (Paul Poenisch)
>State: open
>Class: hw-bug
>Submitter-Id: MUSE
>Arrival-Date: Mon May 1 15:40:01 1995
>Originator: Paul Poenisch
>Organization: MicroUnity Systems Engineering, Inc.
>Release: design rules rev. 5.0
>Environment: Design rules rev. 5.0
Calliope-0, Calliope-1, Orchis-1, as designed Euterpe-1 and Mnemo-1 and Proteus library
>Description: Crack protection ring is designed with 0.5 um by X um slots to meet old 2.0 um maximum metal width rule.
Needs to be redesigned with basket weave pattern.
>How-To-Repeat: NA
>Fix:
>Audit-Trail:
>Unformatted:

From: paulp
Sent: Monday, May 01, 1995 5:40 PM
To: paulp
Cc: geert; hopper; tbr; fwo
Subject: metalcells/2264: Seal ring waffle pattern change

>Number: 2264
>Category: metalcells
>Synopsis: Seal ring waffle pattern change
>Confidential: no
>Severity: critical
>Priority: high
>Responsible: paulp (Paul Poenisch)
>State: open
>Class: hw-bug
>Submitter-Id: MUSE
>Arrival-Date: Mon May 1 15:40:19 1995
>Originator: Paul Poenisch
>Organization: MicroUnity Systems Engineering, Inc.
>Release: design rules rev. 5.0
>Environment: Calliope-0, Calliope-1, Orchis-1, currently designed Euterpe-1 and Mnemo-1 and Proteus library
>Description: Current seal ring design uses 0.5 um x X um slots to meet old 2.0 um maximum metal feature width.
Design needs to be changed to basket weave pattern with 1.5 um x 4.5 um by 1.5 um or 2.0 um x 6.0 um by 2.0 um dimentions.

Cells effected (in Euterpe): padiode, padbl, padbr, padtr and padtl in proteus
>How-To-Repeat: NA
>Fix:
>Audit-Trail:
>Unformatted:

From: hopper (Mark Hofmann)
Sent: Monday, May 01, 1995 2:33 PM
To: Frank Paturzo
Cc: wampler (Kurt Wampler); tau; vant; sysadm
Subject: Re: trex

Frank Paturzo writes:

Odd. I just tested it again and it worked.

How did you test it? (I didn't get a page...)

I also checked the log file and no errors were reported.

Does the log indicate that a page was initiated?

Looks like I'm going to have to pick this apart.

It's a real mess. I'm tempted to rewrite it.

Thanks, Frank. It's important that we have a reliable paging scheme when a critical machine goes down... especially as we approach chip tapeout.

=====

Frank Paturzo

fgp@microunity.com

-hopper

From: pauip (Paul Poenisch)
Sent: Monday, May 01, 1995 10:40 AM
To: Tim B. Robinson
Cc: geert (Geert Rossee)
Subject: Re: processing of Castor/Pollux and Calliope1

Tim Robinson wrote on April 28:

>
>
> Paul Poenisch wrote (on Fri Apr 21):

>
> Calliope-1

>
> The data used to generate this reticle set is not quite as upto date as Orchis
> (which is also not up to current standards) but it should be usable. Unfortun-
> ately when we processed the first lots of this device through the middle
> metal layers (metal 2 through 3) we discovered that the reticle vendor that
> made this set failed to hold the dimentional tolerance needed for this
> process. As a result we are having problems processing this device.

>
> We believe that we can use the current reticle set to produce some Calliope-1's
> but the photomasking engineers will have to hand carry the lots through the
> metal layers. This will result in some slowing of the lots when compared to
> Orchis (which does not have this problem) but should still allow us to get
> enough devices out to allow debug of the device.

>
> Paul, this seems a little different from what was said in the meeting
> last week. If the problem is just poor processing of the masks, then
> why would we expect that necessarily to carry over to mnemo an
> euterpe. (You may remember the issue was Al saying he would not
> process these if they had the same problems.) I realize there are
> other reasons we have to change, but I would like to be sure I fully
> understand the reasoning.

>
> Tim

>
>
Yes, these are two separate problems. The reticle problem effects only
Calliope-1 and should not be a factor in Euterpe or Mnemo (assuming the mask shops get
their act together, and we won't accept the reticles otherwise now that we know what to
look for). This problem effects us in that the photo engineers have to do the processing
for these lots through the metal masking steps. The result is that the lots will proceed
through the line slower than devices that don't have the problem and there will be a
significant impact on other work the engineers need to do to keep the line running and
improve the process to the point that we will start getting yield.

The problems that I talked about last week are separate and effect all the devices (to
varying levels) produced to date and would effect Euterpe and Mnemo if no changes are made
in the Proteus library. These problems have a differnent effect, the line yield for
devices with these problems will be very low. I would expect that we will loose two out
of three wafers that we start on devices that have these problems, additionally only a
portion of each wafer that we do complete will have a chance of yielding due to lift
problems (I would estimate less than half the wafer on average). This means that only
about 15% of the die that we start on these devices will have a chance of yielding.
Additionally wafers that are lost in process will likely be lost late in the flow, via 23
or later, due to a combination of factors.
This is the most expensive and, schedule wise, the most disruptive place to lose wafers.
As an example the lead two lots of Calliope-1 that were suppose to be the first lots to
check for yield (due out last week) were lost at via 34. The follow up lots are due out
this week, so we lost (best case) 3 to
4 days on the schedule sofar.

There were some problems with GNATS which Loretta has now fixed. I have started entering the problems into the system, you should start seeing the reports this morning.

Paul.